Cisco MDS 9000 Family CLI Configuration Guide, Release 3.x

Cisco MDS SAN-OS for Release 3.0(1) Through 3.2(2c)
November 2007

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

Text Part Number: OL-8222-08

Cisco MDS 9000 Family CLI Configuration Guide

CONTENTS

New and Changed Information

This document provides release-specific information for each new and changed feature in the Cisco MDS SAN-OS Release 3.x software. The *Cisco MDS 9000 Family CLI Configuration Guide* is updated to address each new and changed feature in the Cisco MDS SAN-OS Release 3.x software. The latest version of this document is available at the following Cisco Systems website:
http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html

Tip The configuration guides created for earlier releases are also listed at the aforementioned website. Each guide addresses the features introduced in or available in those releases. Select and view the configuration guide that pertains to the software installed in your switch.

To check for additional information about Cisco MDS SAN-OS Release 3.x, refer to the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:
http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_release_notes_list.html

Table 1 summarizes the new and changed features for the *Cisco MDS 9000 Family CLI Configuration Guide, Release 3.x*, and tells you where they are documented. The table includes a brief description of each new feature and the release in which the change occurred.

Table 1 New and Changed Features for Release 3.x

Feature	Description	Changed in Release	Where Documented
Configuring iSCSI	Added the enable and disable command for modules in the iSCSI feature.	3.2(2c)	Chapter 42, “Configuring iSCSI”
Obtaining and Installing Licenses	Removed the prefix “VDH=” from HostID serial number	3.2(2c)	Chapter 3, “Obtaining and Installing Licenses”
N Port Virtualization	N port virtualization educes the number of Fibre Channel domain IDs in SANs.	3.2(1)	Chapter 13, “Configuring N Port Virtualization”
Cisco MDS 18/4-port Multiservice module	Added updates about support or non-support of the module throughout the book.	3.2(1)	Chapter 7, “Software Images” Chapter 14, “Configuring Generation 2 Switches and Modules”

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
Cisco MDS 9222i Multiservice modular switch	Added updates about support or non-support of the switch throughout the book.	3.2(1)	Chapter 7, "Software Images" Chapter 14, "Configuring Generation 2 Switches and Modules"
Cisco MDS 9134 Multilayer Fabric Switch	Added updates about support or non-support of the switch throughout the book.	3.2(1)	Chapter 3, "Obtaining and Installing Licenses" Chapter 7, "Software Images" Chapter 14, "Configuring Generation 2 Switches and Modules"
SANTap enhancements	Provides for 32-bit support and dynamic LUNs.	3.2(1)	Chapter 49, "Configuring SANTap"
CFS enhancements	Support for CFS regions.	3.2(1)	Chapter 6, "Using the CFS Infrastructure"
TACACS+ password expiry notification	Password aging notification is initiated when the TACACS+ server authenticates access to the switch through Telnet or SSH	3.2(1)	Chapter 33, "Configuring RADIUS and TACACS+"
Cisco 32-port switching module	Added a new guideline for copying a saved configuration that contains the **no system default switchport shutdown** command, and the effect that has on E-port configuration.	3.1(3)	Chapter 16, "Configuring PortChannels"
Cisco CompactFlash CRC Checksum Test	Enables users to run the CompactFlash CRC Checksum test and update CompactFlash firmware.	3.1(3)	Chapter 10, "Managing System Hardware"
System default port mode F	Added information about the **system default switchport mode F** feature and command.	3.1(3)	Chapter 5, "Initial Configuration" Chapter 12, "Configuring Interfaces"
Cisco Fabric Switch for HP c-Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter	Added updates about support or non-support of the switches throughout the book.	3.1(2)	Chapter 4, "On-Demand Port Activation Licensing" Chapter 7, "Software Images"
On-Demand Port Activation Licensing	Added port naming conventions f and switch behavior of Cisco Fabric Switch for HP c-Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter.	3.1(2)	Chapter 4, "On-Demand Port Activation Licensing"
Running the CompactFlash Report	Enables users to run the CompactFlash Check utility to generate a report that shows the status of CompactFlash on certain line cards.	3.1(2)	Chapter 10, "Managing System Hardware"
SAN device virtualization	Allows you to create virtual devices that represent physical end devices when configuring switches.	3.1(2)	Chapter 20, "SAN Device Virtualization"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
Enable/disable up/down link	Allows users to enable or disable SNMP link-state traps on specific interfaces.	3.1(2)	Chapter 32, "Configuring SNMP"
Cisco MDS 9124 Fabric Switch support	Updates throughout the book to reflect support of the Cisco MDS 9124 Fabric Switch.	3.1(1)	Chapter 1, "Product Overview"
On-demand port activation licensing	Allows users to buy additional licenses for ports in the Cisco MDS 9124 Fabric Switch, and to also move licenses among ports.	3.1(1)	Chapter 4, "On-Demand Port Activation Licensing"
Non-disruptive upgrades on the Cisco MDS 9124 Fabric Switch	Describes non-disruptive upgrades on the Cisco MDS 9124 Fabric Switch.	3.1(1)	Chapter 7, "Software Images"
Removal of restrictions on oversubscription ratios	Allows users to remove any restrictions on maximum oversubscription ratios.	3.1(1)	Chapter 14, "Configuring Generation 2 Switches and Modules"
FICON Tape Acceleartion	Provides acceleration for FICON tape write operations over FCIP for the IBM Virtual Tape Server (VTS) and tape libraries that support the 3490 command set.	3.1(1)	Chapter 28, "Configuring FICON"
IP Network Simulator	Allows users to simulate network conditions to test the impact of network latency for FCIP or iSCSI.	3.1(1)	Chapter 58, "Troubleshooting Your Fabric"
Generation 2 switching modules	Provides default support for Fibre Channel ACL adjacency sharing on Generation 2 switching modules.	3.0(3)	Chapter 14, "Configuring Generation 2 Switches and Modules"
Command scheduler remote user support	Allows remote users to configure command scheduler jobs.	3.0(3)	Chapter 18, "Scheduling Maintenance Jobs"
IVR zones and zone members	Increases the limits for IVR zones to 8000 and for IVR zone members to 20,000	3.0(3)	Chapter 22, "Configuring Inter-VSAN Routing"
RLIR messages	Allows you to specify a server to receive Registered Link Incident Report (RLIR) frames.	3.0(3)	Chapter 28, "Configuring FICON"
User configuration limit	Sets the maximum number of users on a switch to 256.	3.0(3)	Chapter 37, "Configuring Users and Common Roles"
show tech-support command	Allows you to specify new options for the **show tech-support** command for specific features.	3.0(3)	Chapter 58, "Troubleshooting Your Fabric"
install ssi command	Copies the SSI boot image file to the SSM modflash:.	3.0(2)	Chapter 11, "Managing Modules"
Domain manager fast restart	Allows the domain manager to quickly recover from a principal link failure when a backup link is available.	3.0(2)	Chapter 17, "Configuring Domain Parameters"
FICON port prohibiting default	Allows you to change the default setting for port prohibiting.	3.0(2)	Chapter 28, "Configuring FICON"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
CLI enhancements	Includes the following CLI enhancements: • User-defined command variables. • User-defined aliases for common commands. • The **pwc** command displays the list of commands entered to reach a command submode. • Command variable support in the **run-script** command.	3.0(1)	Chapter 2, "Before You Begin"
Configuration format change	Describes the multiple entry format for displaying interface configuration information in the **show running-config** and **show startup-config** command outputs.	3.0(1)	Chapter 2, "Before You Begin" Chapter 12, "Configuring Interfaces"
Supervisor-2 module support	Includes support for the following Supervisor-2 module features: • Configuring modem parameters on the console port and COM1 port. • Allowing 1000 Mbps speed on the management port.	3.0(1)	Chapter 5, "Initial Configuration"
CFS over IP	Allows CFS distributions over IP connections.	3.0(1)	Chapter 6, "Using the CFS Infrastructure"
Configuration check	Describes the changes to the **show incompatibility system** command that indicate the commands to use to disable features before downgrading to an earlier release of the system image.	3.0(1)	Chapter 7, "Software Images"
Supervisor module management procedures	Includes the following recommended supervisor module management procedures: • Preparing to remove supervisor modules from Cisco MDS 9500 Series Directors containing both Generation 1 and Generation 2 switching modules. • Migrating from Supervisor-1 modules to Supervisor-2 modules in the Cisco MDS 9500 Series Directors.	3.0(1)	Chapter 7, "Software Images"
boot auto-copy command enabled by default	Changes the default state for the **boot auto-copy** command to enabled.	3.0(1)	Chapter 9, "Configuring High Availability"
Crossbar removal procedures	Provides the recommended procedures to prepare to remove crossbars from Cisco MDS 9500 Series Directors.	3.0(1)	Chapter 10, "Managing System Hardware"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
N-port identifier virtualization (NPIV)	Provides support for multiple N port identifiers to support multiple applications on a server accessing an MDS switch.	3.0(1)	Chapter 12, "Configuring Interfaces"
Auto port mode support change	Disallows configuring auto port mode on Storage Services Module (SSM) Fibre Channel interfaces.	3.0(1)	Chapter 12, "Configuring Interfaces"
Generation 2 switching module support	Describes how to configure interfaces on the Generation 2 Fibre Channel switching modules.	3.0(1)	Chapter 14, "Configuring Generation 2 Switches and Modules"
SFP diagnostic information	Describes the **show interface transceiver** command change to display SFP diagnostic information on Generation 2 modules.	3.0(1)	Chapter 14, "Configuring Generation 2 Switches and Modules"
CFS support for allowed domain ID lists	Allows the allowed domain ID lists to be distributed in the fabric using the CFS infrastructure.	3.0(1)	Chapter 17, "Configuring Domain Parameters"
IVR enhancements	Includes the following IVR enhancements: • Allowing more than one active IVR service group. • Renaming IVR zones (IVZs). • Renaming IVR zone sets (IVZSs). • Copying the active IVZS to the full IVZS to be edited and reactivated. • Copying the active IVR topology to the manually configured IVR topology. • Copying the active IVR service group database to the configured IVR service group database. • Clearing the configured IVR service group database.	3.0(1)	Chapter 22, "Configuring Inter-VSAN Routing"
Increased zone limit per VSAN	Increases the maximum number of zones per VSAN from 2000 to 8000.	3.0(1)	Chapter 23, "Configuring and Managing Zones"
Zone analysis	Provides a means to analyze zone characteristics using the **show zone analysis** command.	3.0(1)	Chapter 23, "Configuring and Managing Zones"
Device alias rename	Allows existing device aliases to be renamed.	3.0(1)	Chapter 24, "Distributing Device Alias Services"
In-order-delivery enhancement	Ensures that frames are delivered in order within the switch latency drop period.	3.0(1)	Chapter 25, "Configuring Fibre Channel Routing Services and Protocols"
CFS support for RCSN	Allows the RSCN timer value to be distributed in the fabric using the CFS infrastructure.	3.0(1)	Chapter 26, "Managing FLOGI, Name Server, FDMI, and RSCN Databases"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
RSCN timer configuration	Allows the RSCN timer value to be configured.	3.0(1)	Chapter 26, "Managing FLOGI, Name Server, FDMI, and RSCN Databases"
FICON port numbering	Provides information on the changed default port numbering scheme for Generation 2 hardware and how to assign port numbers when a switch has more than 255 ports.	3.0(1)	Chapter 28, "Configuring FICON"
fcid-last-byte command deprecated	Does not support the **fcid-last-byte** command.	3.0(1)	Chapter 28, "Configuring FICON"
FICON port swapping	Provides the ability to port swap using the interface identifier when there are duplicate port numbers on a switch.	3.0(1)	Chapter 28, "Configuring FICON"
SSH authentication enhancement	Provides digital certificate support for host authentication.	3.0(1)	Chapter 37, "Configuring Users and Common Roles"
AAA server enhancements	Includes the following AAA server enhancements: • Monitoring and validating the availability of remote AAA servers. • Allowing users to specify a remote AAA server name at login. • Displaying AAA server statistics.	3.0(1)	Chapter 33, "Configuring RADIUS and TACACS+"
MSCHAP	Provides support for the Microsoft Challenge Handshake Authentication Protocol (MSCHAP).	3.0(1)	Chapter 33, "Configuring RADIUS and TACACS+"
Change to **show ip access-list** command	Deprecates the **usage** option.	3.0(1)	Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists"
IPv6 access control lists (IPv6-ACLs)	Describes the support for IPv6-ACLs.	3.0(1)	Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists"
Certificate authorities and digital certificates	Describes interoperating with certificate authorities and using digital certificates for secure communication with peers.	3.0(1)	Chapter 35, "Configuring Certificate Authorities and Digital Certificates"
IKE digital certificate support	Allows IKE to use digital certificates for authentication instead of using preshared keys.	3.0(1)	Chapter 36, "Configuring IPsec Network Security"
IKE fully qualified domain name (FQDN)	Includes using FQDNs, as well as IPv4 addresses, for the following IKE features: • Preshared keys • Identity mode	3.0(1)	Chapter 36, "Configuring IPsec Network Security"
Fabric binding for Fibre Channel	Supports fabric binding for Fibre Channel VSANs as well as FICON VSANs.	3.0(1)	Chapter 39, "Configuring Fabric Binding"
FCIP read tape acceleration	Supports tape read acceleration over FCIP interfaces as well as tape write acceleration.	3.0(1)	Chapter 40, "Configuring FCIP"

Table 1 New and Changed Features for Release 3.x (continued)

Feature	Description	Changed in Release	Where Documented
SAN extension tuner enhancement	Describes how to assign tape read and write commands to N ports.	3.0(1)	Chapter 41, "Configuring the SAN Extension Tuner"
iSCSI server load balancing (iSLB)	Provides information on how to easily configure large iSCSI deployments.	3.0(1)	Chapter 42, "Configuring iSCSI"
iSNS cloud discovery	Provides information to iSNS on the various interfaces reachable from an initiator by partitioning the interfaces on a switch into disjointed IP clouds.	3.0(1)	Chapter 42, "Configuring iSCSI"
Dynamic initiator modes	Allows configuration of dynamic initiator modes iSCSI, iSLB, and deny log in to the MDS switch.	3.0(1)	Chapter 42, "Configuring iSCSI"
IPv6	Provides support for IP version 6 (IPv6).	3.0(1)	Chapter 43, "Configuring IP Services" Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces"
RMON high capacity alarms	Provides the **show rmon high capacity alarms** command to display RMON high capacity alarm values.	3.0(1)	Chapter 51, "Configuring RMON"
Call Home enhancement	Allows customization of alert group messages.	3.0(1)	Chapter 54, "Configuring Call Home"
QoS behavior	Provides information about the behavior of QoS with different combinations of Generation 1 and Generation 2 switching modules.	3.0(1)	Chapter 56, "Configuring Fabric Congestion Control and QoS"
On-line system health maintenance (OHMS) enhancements	Includes the following OHMS enhancements: • Configuring the global frame length for loopback test for all modules on the switch. • Specifying frame count and frame length on for the loopback test on a specific module. • Configuring source and destination ports for external loopback tests. • Providing serdes loopback test to check hardware.	3.0(1)	Chapter 59, "Monitoring System Processes and Logs"
On-board failure logging (OBFL)	Describes OBFL, how to configure it for Generation 2 modules, and how to display the log information.	3.0(1)	Chapter 59, "Monitoring System Processes and Logs"

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family CLI Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Product Overview	Presents an overview of the Cisco MDS 9000 Family of multilayer switches and directors.
Chapter 2	Before You Begin	Describes the command-line interface (CLI).
Chapter 3	Obtaining and Installing Licenses	Describes license types, procedure, installation, and management for the Cisco MDS SAN-OS software.
Chapter 4	On-Demand Port Activation Licensing	Describes how to configure and manage on-demand ports for switches that support on-demand port activation licensing.
Chapter 5	Initial Configuration	Provides initial switch configuration options and switch access information.
Chapter 6	Using the CFS Infrastructure	Explains the use of the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution.
Chapter 7	Software Images	Describes how to install and upgrade software images
Chapter 8	Working with Configuration Files	Describes the initial configuration of the switches using the configuration files so they can be accessed by other devices

Chapter	Title	Description
Chapter 9	Configuring High Availability	Describes the high availability feature including switchover mechanisms.
Chapter 10	Managing System Hardware	Explains switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information.
Chapter 11	Managing Modules	Explains how to display and analyze the status of each module and specifies the power on and power off process for modules.
Chapter 12	Configuring Interfaces	Explains Generation 1 and Generation 2 module port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
Chapter 13	Configuring N Port Virtualization	Explains how to configure NPV devices to reduce excessive Fibre Channel domain IDs in SANs.
Chapter 14	Configuring Generation 2 Switches and Modules	Explains configuration concepts for Generation 2 module ports and interfaces.
Chapter 15	Configuring Trunking	Explains TE ports and trunking concepts.
Chapter 16	Configuring PortChannels	Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.
Chapter 17	Configuring Domain Parameters	Explains the Fibre Channel domain (fcdomain) feature, which includes principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions.
Chapter 18	Scheduling Maintenance Jobs	Describes the Cisco MDS command scheduler feature that helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family.
Chapter 19	Configuring and Managing VSANs	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs.
Chapter 20	SAN Device Virtualization	Describes how to configure virtual devices to represent physical end devices for switches running Cisco MDS SAN-OS Release 3.1(2) and later.
Chapter 21	Creating Dynamic VSANs	Defines the Dynamic Port VSAN Membership (DPVM) feature that is used to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches.

Chapter	Title	Description
Chapter 22	Configuring Inter-VSAN Routing	Provides details on sharing resources across VSANs using the inter-VSAN Routing (IVR) feature.
Chapter 23	Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Chapter 24	Distributing Device Alias Services	Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis.
Chapter 25	Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Chapter 26	Managing FLOGI, Name Server, FDMI, and RSCN Databases	Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases.
Chapter 27	Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Chapter 28	Configuring FICON	Provides details on the FI-bre CON-nection (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS switches.
Chapter 29	Advanced Features and Concepts	Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
Chapter 30	Configuring FIPS	Describes Federal Information Processing Standards (FIPS) implementation in SAN-OS, and how to enable your system to operate in a FIPS-compliant mode.
Chapter 37	Configuring Users and Common Roles	Describes how to configure users and common roles.
Chapter 32	Configuring SNMP	Provides details on how you can use SNMP to modify a role that was created using CLI.
Chapter 33	Configuring RADIUS and TACACS+	Discusses the AAA parameters, user profiles, and RADIUS authentication security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options.
Chapter 34	Configuring IPv4 and IPv6 Access Control Lists	Describes the IPv4 and IPv6 static routing feature and its use to route traffic between VSANs.
Chapter 35	Configuring Certificate Authorities and Digital Certificates	Describes how to interoperate with Certificate Authorities (CAs) and use digital certificates for secure, scalable communication.

Chapter	Title	Description
Chapter 36	Configuring IPsec Network Security	Provides details on the digital certificates, IP Security Protocol (IPsec) open standards, and the Internet Key Exchange (IKE) protocol that it uses to handle protocol and algorithm negotiation.
Chapter 37	Configuring FC-SP and DHCHAP	Describes the DHCHAP protocol, an FC-SP protocol, that provides authentication between Cisco MDS 9000 Family switches and other devices.
Chapter 38	Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.
Chapter 39	Configuring Fabric Binding	Describes the fabric binding security feature for VSANs, which ensures that ISLs are only enabled between specific switches.
Chapter 40	Configuring FCIP	Describes how the switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.
Chapter 41	Configuring the SAN Extension Tuner	Explains the SAN extension tuner (SET) feature that optimizes FCIP performance.
Chapter 42	Configuring iSCSI	Describes the iSCSI feature that is specific to the IPS module and is available in the Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.
Chapter 43	Configuring IP Services	Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.
Chapter 44	Configuring IP Storage	Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together through IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol.
Chapter 45	Configuring IPv4 for Gigabit Ethernet Interfaces	Describes the IPv4 protocol support provided by Cisco MDS 9000 Family switches.
Chapter 46	Configuring IPv6 for Gigabit Ethernet Interfaces	Describes the IPv6 protocol support provided by Cisco MDS 9000 Family switches.
Chapter 47	Configuring SCSI Flow Services and Statistics	Describes the SCSI flow services and SCSI flow statistics, the Intelligent Storage Services.
Chapter 48	Configuring Fibre Channel Write Acceleration	Describes Fibre Channel Write Acceleration support and configuration.
Chapter 49	Configuring SANTap	Describes SANTap support and configuration.
Chapter 50	Configuring NASB	Describes NASB support and configuration.
Chapter 51	Configuring RMON	Provides details on using RMONs to configure alarms and events.

Chapter	Title	Description
Chapter 52	Monitoring Network Traffic Using SPAN	Describes the Switched Port Analyzer (SPAN), SPAN sources, filters, SPAN sessions, SD port characteristics, and configuration details.
Chapter 53	Configuring System Message Logging	Describes how system message logging is configured and displayed.
Chapter 54	Configuring Call Home	Provides details on the Call Home service and includes information on Call Home, event triggers, contact information, destination profiles, and e-mail options.
Chapter 55	Configuring Fabric Configuration Servers	Describes how the fabric configuration server (FCS) feature is configured and displayed.
Chapter 56	Configuring Fabric Congestion Control and QoS	Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features.
Chapter 57	Configuring Port Tracking	Provides information about a port tracking feature that provides a faster recovery from link failures.
Chapter 58	Troubleshooting Your Fabric	Describes basic troubleshooting methods used to resolve issues with switches.
Chapter 59	Monitoring System Processes and Logs	Provides information on displaying system processes and status. It also provides information on configuring core and log files, HA policy, heartbeat and watchdog checks, and upgrade resets.
Appendix A	Configuration Limits for Cisco MDS SAN-OS Release 3.x	Lists the Cisco verified limits and maximum limits for switches running Cisco MDS SAN-OS Release 3.x.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
italic font	Arguments for which you supply values are in italics.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
`italic screen font`	Arguments for which you supply values are in italic screen font.
`< >`	Nonprinting characters, such as passwords, are in angle brackets.

[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

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

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

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents.

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Storage Services Interface Images*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*

Compatibility Information

- *Cisco MDS 9000 SAN-OS Hardware and Software Compatibility Information*
- *Cisco MDS 9000 Family Interoperability Support Matrix*
- *Cisco MDS Storage Services Module Interoperability Support Matrix*
- *Cisco MDS SAN-OS Release Compatibility Matrix for Storage Service Interface Images*

Regulatory Compliance and Safety Information

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*

Hardware Installation

- *Cisco MDS 9124 Multilayer Fabric Switch Quick Start Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*

Cisco Fabric Manager

- *Cisco MDS 9000 Family Fabric Manager Quick Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Family Data Mobility Manager Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager Database Schema*

Command-Line Interface

- *Cisco MDS 9000 Family Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Storage Services Module Software Installation and Upgrade Guide*
- *Cisco MDS 9000 Family CLI Quick Configuration Guide*
- *Cisco MDS 9000 Family CLI Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*

Troubleshooting and Reference

- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family MIB Quick Reference*
- *Cisco MDS 9000 Family SMI-S Programming Reference*
- *Cisco MDS 9000 Family System Messages Reference*

Installation and Configuration Note

- *Cisco MDS 9000 Family SSM Configuration Note*
- *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Note*
- *Cisco 10-Gigabit X2 Transceiver Module Installation Note*
- *Cisco MDS 9000 Family CWDM SFP Installation Note*
- *Cisco MDS 9000 Family CWDM Passive Optical System Installation Note*

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, obtaining support, providing documentation feedback, security guidelines, and also recommended aliases and general Cisco documents, see the monthly *What's New* in Cisco Product Documentation, which also lists all new and revised Cisco technical documentation, at:

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

PART 1

Getting Started

CHAPTER 1

Product Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offers intelligent fabric-switching services that realize maximum performance while ensuring high reliability levels. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This powerful combination enables highly available, scalable storage networks that provide intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

This chapter lists the hardware features for the Cisco MDS 9000 Family and describes its software features. It includes the following sections:

Hardware Overview

This section provides an overview of the following Cisco MDS 9000 Family of multilayer directors and fabric switches:

- Cisco MDS 9500 Series multilayer directors
  - Cisco MDS 9513 multilayer director
  - Cisco MDS 9509 multilayer director
  - Cisco MDS 9506 multilayer director
- Cisco MDS 9200 Series fabric switches
  - Cisco MDS 9222i multilayer fabric switch
  - Cisco MDS 9216i multilayer fabric switch
  - Cisco MDS 9216A multilayer fabric switch
  - Cisco MDS 9216 multilayer fabric switch
- Cisco MDS 9100 Series fixed configuration fabric switches
  - Cisco MDS 9140 multilayer switch
  - Cisco MDS 9134 multilayer switch
  - Cisco MDS 9124 multilayer switch
  - Cisco MDS 9120 multilayer switch

- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

Cisco MDS 9500 Series Multilayer Directors

The Cisco MDS 9500 Series includes the following multilayer, modular directors:

- The Cisco MDS 9513 Director, which has thirteen slots, two of which (slot 7 and slot 8) are reserved for the Supervisor-2 modules, and can accommodate up to eleven hot-pluggable switching or services modules.
- The Cisco MDS 9509 Director, which has nine slots, two of which (slot 5 and slot 6) are reserved for the Supervisor-1 modules or Supervisor-2 modules, and can accommodate up to seven hot-pluggable switching or services modules.
- The Cisco MDS 9506 Director, which has six slots, two of which (slot 5 and slot 6) are reserved for the Supervisor-1 modules or Supervisor-2 modules, and can accommodate up to four hot-pluggable switching or services modules.

Note Supervisor-1 modules and Supervisor-2 modules can only operate in the same chassis during migration. See the "Migrating from Supervisor-1 Modules to Supervisor-2 Modules" section on page 7-33.

The two supervisor modules ensure high availability and traffic load balancing capabilities. The standby supervisor module provides redundancy if the active supervisor module fails. Supervisor-1 modules provide management access through a 10/100BASE-T Ethernet port switch and an RS-232 serial port. Supervisor-2 modules provide management access through a 10/100/1000BASE-T Ethernet port switch and an RS-232 serial port.

Note As of Cisco MDS SAN-OS release 3.2(1), the USB ports on the Supervisor-2 module are supported. USB flash drives connected to these ports may be used for the same functions as media in the external compact flash slot.

The Cisco MDS 9500 Series directors support the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 18/4-port Multiservice module (MSM-18/4)
- 18/4-port Multiservice module FIPS
- 18-port 4-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module

- Storage Services Module (SSM)

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Cisco MDS 9200 Series Fabric Switches

The Cisco MDS 9200 Series includes the following multilayer switches supporting multiprotocol capabilities:

- Cisco MDS 9222i
- Cisco MDS 9216i
- Cisco MDS 9216A
- Cisco MDS 9216

Cisco MDS 9216i Multiprotocol Fabric Switch

The Cisco MDS 9216i multiprotocol fabric switch has two slots, one of which is reserved for the integrated supervisor module and the other for switching or services modules. The supervisor module provides supervisor functions and has 14 standard Fibre Channel ports and two multiprotocol ports that can support FCIP and iSCSI protocols simultaneously.

The Cisco MDS 9200 multilayer fabric switches support the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Storage Services Module (SSM)

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide*.

Cisco MDS 9222i, Cisco MDS 9216A and Cisco MDS 9216 Multilayer Fabric Switches

The Cisco MDS 9222i, Cisco MDS 9216A and Cisco MDS 9216 multilayer fabric switches have two slots, one of which is reserved for the integrated supervisor module and the other for a switching or services module. The supervisor module provides supervisor functions and has 16 standard Fibre Channel ports.

The Cisco MDS 9222i multilayer fabric switch supports the following switching and services modules:

- 12-port, 24-port, and 48-port 4-Gbps Fibre Channel switching modules
- 4-port 10-Gbps Fibre Channel switching module
- 18/4-port Multiservice Module

- 18/4-port Multiservice FIPS Module with Federal Information Processing Standard (FIPS) 140-2 Level-3 validation
- 32-port Storage Services Module
- 8-port IP Storage Services Module

The Cisco MDS 9216A multilayer fabric switch supports the following switching and services modules:

- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module

The Cisco MDS 9216 multilayer fabric switch supports the following switching and services modules:

- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 14/2-port Multiprotocol Services (MPS-14/2) module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* and the *Cisco MDS 9216 Switch Hardware Installation Guide*.

Cisco MDS 9100 Series Fixed Configuration Fabric Switches

Cisco MDS 9100 Series includes the following multilayer, fixed configuration (non-modular) switches:

- Cisco MDS 9140 with 40 ports (8 full-rate ports, 32 host-optimized ports)
- Cisco MDS 9134 with 34 ports (24-port base with 8-port license for growth; two 10 Gbps ports can be activated independently in 24-port or 32-port configurations)
  - On-demand port activation licensing
  - Non-disruptive upgrades
- Cisco MDS 9124 with 24 ports (8-port base with 8-port license for growth)

  Also includes:
  - On-demand port activation licensing

  - Non-disruptive upgrades
- Cisco MDS 9120 with 20 ports (4 full-rate ports, 16 host-optimized ports)
- Cisco Fabric Switch for HP c-Class BladeSystem (24 ports; 14 internal 2/4 Gbps, and 6 full-rate ports)
- Cisco Fabric Switch for IBM BladeCenter (20 ports; 14 internal 2/4 Gbps, and 6 external full-rate ports)

These fixed configuration switches are packaged in 1 RU enclosures and provide 1-Gbps, 2-Gbps, 4-Gbps, or 10 Gbps autosensing Fibre Channel ports. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access.

Note Switches in the Cisco MDS 9100 Series do not have a COM1 port (RS-232 serial port).

Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Cisco SAN-OS Software Configuration

This section describes the tools you can use to configure SAN-OS software, and provides an overview of the software configuration process with links to the appropriate chapters.

This section includes the following topics:

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs (see Figure 1-1).

- The command-line interface (CLI) can manage Cisco MDS 9000 Family switches using Telnet, SSH, or a serial connection.
- The Cisco MDS 9000 Fabric Manager, a Java-based graphical user interface, can manage Cisco MDS 9000 Family switches using SNMP.

Figure 1-1 Tools for Configuring Cisco SAN-OS Software

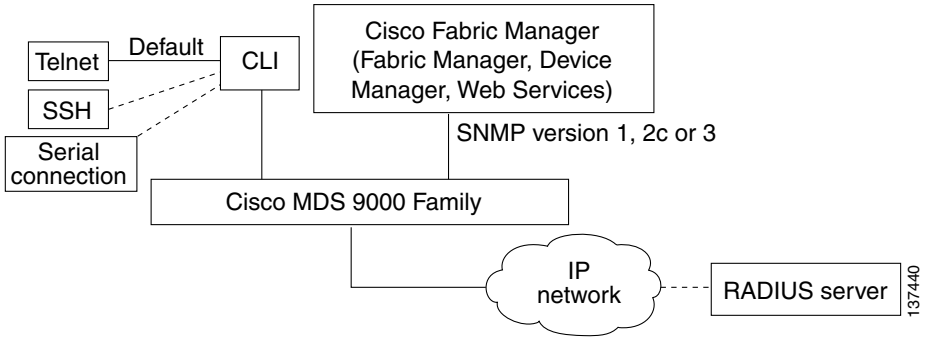

CLI

With the CLI, you can type commands at the switch prompt, and the commands are executed when you press the **Enter** key. The CLI parser provides command help, command completion, and keyboard sequences that allow you to access previously executed commands from the buffer history.

Continue reading this document for more information on configuring the Cisco MDS switch using the CLI.

Cisco MDS 9000 Fabric Manager

The Cisco Fabric Manager is a set of network management tools that supports Secure Simple Network Management Protocol version 3 (SNMPv3) and legacy versions. The Cisco Fabric Manager applications are:

- Fabric Manager Client—provides a graphical user interface (GUI) that displays real-time views of your network fabric, and lets you manage the configuration of Cisco MDS 9000 Family devices and third-party switches.
- Fabric Manager Server—performs advanced monitoring, troubleshooting, and configuration for multiple fabrics. It must be started before running the Fabric Manager Client. It can be accessed by up to 16 Fabric Manager Clients at a time.
- Device Manager—presents two views of a switch.
  - Device View displays a continuously updated physical representation of the switch configuration, and provides access to statistics and configuration information for a single switch.
  - Summary View presents real-time performance statistics of all active interfaces and channels on the switch for Fibre Channel and IP connections.
- Fabric Manager Web Services—allows operators to monitor MDS events, performance, and inventory, and perform minor configuration tasks from a remote location using a web browser.
- Performance Manager—provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts that can be viewed with any web browser using Fabric Manager Web Services.

The Cisco Fabric Manager applications are an alternative to the CLI for most switch configuration commands.

Note Resource Manager Essentials (RME) versions 3.4 and 3.5 provide support for switches in the Cisco MDS 9000 Family. Device Updates (DU) are available on Cisco.com (http://www.cisco.com/).

For more information on configuring the Cisco MDS switch using the Cisco MDS 9000 Family Fabric Manager, refer to the *Cisco MDS 9000 Fabric Manager Configuration Guide*.

Software Configuration Overview

This section provides an overview of the Cisco SAN-OS configuration process and includes the following topics:

Basic Configuration

These sections contain the minimum information you need to get your switch up and running.

- Preparing to configure the switch (Chapter 2, “Before You Begin”)
- Installing licenses (Chapter 3, “Obtaining and Installing Licenses”)
- Activating additional ports (Chapter 4, “On-Demand Port Activation Licensing”)
- Configuring the minimum requirements:
  - Initial configuration (Chapter 5, “Initial Configuration”)
  - VSANs (Chapter 19, “Configuring and Managing VSANs.”)
  - Interfaces (Chapter 12, “Configuring Interfaces”)
  - Zones and zone sets (Chapter 23, “Configuring and Managing Zones.”)

Advanced Configuration

These sections contain additional configuration information for SAN-OS software and the MDS 9000 Family of switches and includes the following topics:

Switch Configuration

- On-demand port activation licensing (Chapter 4, “On-Demand Port Activation Licensing”)
- N Port virtualization (Chapter 13, “Configuring N Port Virtualization”)
- Generation 2 switching modules (Chapter 14, “Configuring Generation 2 Switches and Modules”)
- High Availability (Chapter 9, “Configuring High Availability”)
- Trunking (Chapter 15, “Configuring Trunking”)
- PortChannels (Chapter 16, “Configuring PortChannels”)
- Domains (Chapter 17, “Configuring Domain Parameters”)
- Schedule maintenance jobs (Chapter 18, “Scheduling Maintenance Jobs”)

Fabric Configuration

- Dynamic VSANs (Chapter 21, “Creating Dynamic VSANs”)
- SAN device virtualization (Chapter 20, “SAN Device Virtualization”)
- Inter-VSAN Routing (Chapter 22, “Configuring Inter-VSAN Routing”)
- Device alias distribution (Chapter 24, “Distributing Device Alias Services”)
- FSPF (Chapter 25, “Configuring Fibre Channel Routing Services and Protocols”)

- FLOGI (Chapter 26, "Managing FLOGI, Name Server, FDMI, and RSCN Databases")
- SCSI (Chapter 27, "Discovering SCSI Targets")
- FICON (Chapter 28, "Configuring FICON")
- Switch interoperability (Chapter 29, "Advanced Features and Concepts")

Security

- Users and Roles (Chapter 37, "Configuring Users and Common Roles")
- SNMP (Chapter 32, "Configuring SNMP")
- RADIUS and TACACS+ (Chapter 33, "Configuring RADIUS and TACACS+")
- Access lists for IPv4 and IPv6 (Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists")
- Digital certificates (Chapter 35, "Configuring Certificate Authorities and Digital Certificates")
- IPsec for network security (Chapter 36, "Configuring IPsec Network Security")
- FC-SP for fabric security (Chapter 37, "Configuring FC-SP and DHCHAP")
- Port security (Chapter 38, "Configuring Port Security")
- Fabric binding (Chapter 39, "Configuring Fabric Binding")

IP Services

- FCIP (Chapter 40, "Configuring FCIP")
- SAN extension tuner (Chapter 41, "Configuring the SAN Extension Tuner")
- iSCSI (Chapter 42, "Configuring iSCSI")
- IP services (Chapter 43, "Configuring IP Services")
- IP storage (Chapter 44, "Configuring IP Storage")
- IPv4 (Chapter 45, "Configuring IPv4 for Gigabit Ethernet Interfaces"
- IPv6 (Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces")

Intelligent Storage Services

- SCSI flow services (Chapter 47, "Configuring SCSI Flow Services and Statistics")
- Fibre Channel write acceleration (Chapter 48, "Configuring Fibre Channel Write Acceleration"
- SANTap (Chapter 49, "Configuring SANTap")
- NASB (Chapter 50, "Configuring NASB")

Network and Switch Monitoring

- RMON (Chapter 51, "Configuring RMON")
- SPAN (Chapter 52, "Monitoring Network Traffic Using SPAN")
- System message logging (Chapter 53, "Configuring System Message Logging")
- Call Home (Chapter 54, "Configuring Call Home")
- Fabric configuration servers (Chapter 55, "Configuring Fabric Configuration Servers")

Traffic Management

- QoS (Chapter 56, “Configuring Fabric Congestion Control and QoS”)
- Port tracking (Chapter 57, “Configuring Port Tracking”)

CHAPTER 2

Before You Begin

This chapter prepares you to configure switches from the CLI. It also describes the CLI command modes and how to use the switch file systems.

This chapter includes the following sections:

About the Switch Prompt

Note Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for installation and connection instructions.

Once the switch is powered on successfully, you see the default switch prompt (`switch#`) as shown in Example 2-1.

Example 2-1 Output When Switch Boots Up

```
Auto booting bootflash:/boot-279 bootflash:/system_image;...
Booting kickstart image:bootflash:/boot-279....
.........................................Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
Uncompressing system image: bootflash:/system_image
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3

<<<<<<SAN OS bootup log messages>>>>>>

         ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Use ctrl-c to abort configuration dialog at any prompt.

Basic management setup configures only enough connectivity for
management of the system.

Would you like to enter the basic configuration dialog (yes/no): yes

<<<<<<after configuration>>>>>>

switch login:admin101
Password:*******
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2006, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software may be covered under the GNU Public
License or the GNU Lesser General Public License. A copy of
each such license is available at
http://www.gnu.org/licenses/gpl.html and
http://www.gnu.org/licenses/lgpl.html
switch#
```

You can perform embedded CLI operations, access command history, and use command parsing functions at this prompt. The switch gathers the command string upon detecting an **Enter** (CR) and accepts commands from a terminal.

Default Switch Roles

By default, two roles exist in all switches:

- Network operator—Has permission to view the configuration.
- Network administrator—Has permission to perform all functions and to set up to 64 permission levels based on user roles and groups.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have the correct permission as specified in the description of the command. See Chapter 37, "Configuring Users and Common Roles."

Using the CLI

This section includes the following topics:

CLI Command Modes

Switches in the Cisco MDS 9000 Family have two main command modes—user EXEC mode and configuration mode. The commands available to you depend on the mode you are in. To obtain a list of available commands in either mode, type a question mark (?) at the system prompt.

Table 2-1 lists and describes the two commonly used modes, how to enter the modes, and the resulting system prompts. The system prompt helps you identify which mode you are in and hence, which commands are available to you.

Table 2-1 Frequently Used Switch Command Modes

Mode	Description of Use	How to Access	Prompt
EXEC	Enables you to temporarily change terminal settings, perform basic tests, and display system information. **Note** Changes made in this mode are generally not saved across system resets.	At the switch prompt, enter the required EXEC mode command.	`switch#`
Configuration mode	Enables you to configure features that affect the system as a whole. **Note** Changes made in this mode are saved across system resets if you save your configuration. See the "Saving a Configuration" section on page 2-14.	From EXEC mode, enter the **config terminal** command.	`switch(config)#`

You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **config terminal** command to **conf t**.

Note Do not enter percent (%), pound (#), ellipsis (...), vertical bar (|), less than or great than (< >), brackets ([ ]), or braces ({ }) in command lines. These characters have special meaning in Cisco SAN-OS text strings.

CLI Command Hierarchy

The CLI commands are organized hierarchically, with commands that perform similar functions grouped under the same level. For example, all commands that display information about the system, configuration, or hardware are grouped under the **show** command, and all commands that allow you to configure the switch are grouped under the **config terminal** command. Figure 2-1 illustrates a portion of the **config terminal** command hierarchy.

Figure 2-1 CLI Command Hierarchy Example

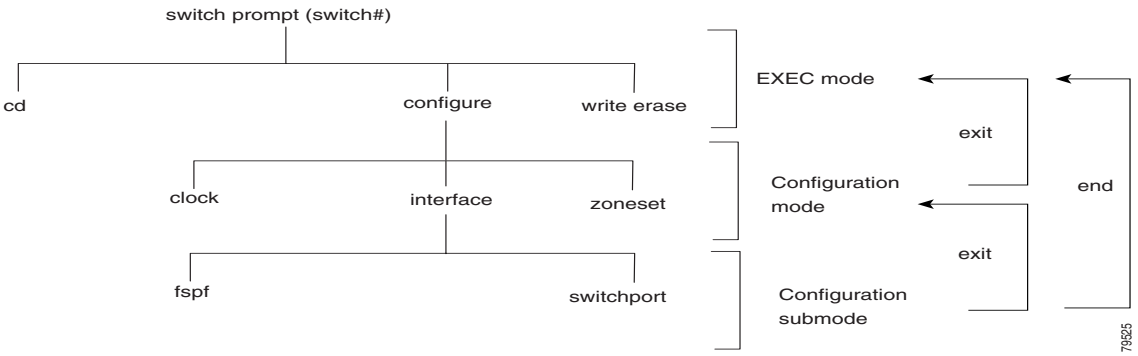

To start executing commands, enter the command at the top level of the hierarchy (EXEC mode). For example, to configure a Fibre Channel interface, use the **config terminal** command. Once you are in configuration mode, issue the **interface** command. When you are in the interface configuration submode, you can query the available commands there.

The following example shows how to query the available commands in the interface submode:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# interface fc1/1
switch(config-if)# ?
Interface configuration commands:
  channel-group  Add to/remove from a port-channel
  do             EXEC command
  exit           Exit from this submode
  fcdomain       Configure fcdomain parameters
  fspf           Configure FSPF parameters
  no             Negate a command or set its defaults
  rspan-tunnel   Configure remote span tunnel interface
  shutdown       Enable/disable an interface
  switchport     Configure switchport parameters
```

EXEC Mode Options

When you start a session on the switch, you begin in EXEC mode. Based on the role or group to which you belong, you have access to limited commands or to all commands (see the "Role-Based Authorization" section on page 37-1). From EXEC mode, you can enter configuration mode. Most of the EXEC commands are one-time commands, such as **show** commands, which display the current configuration status. Here is a list of EXEC mode commands:

```
switch# ?
Exec commands:
  attach        Connect to a specific linecard
  callhome      Callhome commands
  cd            Change current directory
  clear         Reset functions
  cli           CLI commands
  clock         Manage the system clock
  config        Enter configuration mode
  copy          Copy from one file to another
  debug         Debugging functions
  delete        Delete a file
  dir           List files in a directory
  discover      Discover information
  exit          Exit from the EXEC
  fcping        Ping an N-Port
  fctrace       Trace the route for an N-Port.
  find          Find a file below the current directory
  format        Format disks
  gunzip        Uncompresses LZ77 coded files
  gzip          Compresses file using LZ77 coding
  install       Upgrade software
  ips           Various sibyte module related commands
  isns          Re-registers specified interface with isns server
  mkdir         Create new directory
  modem         Modem commands
  move          Move files
  nasb          NASB control functionality
  no            Disable debugging functions
  ping          Send echo messages
  port-channel  Port-Channel related commands
```

```
  purge          Deletes unused data
  pwd            View current directory
  reload         Reboot the entire box
  rmdir          Delete a directory
  run-script     Run shell scripts
  send           Send message to open sessions
  setup          Run the basic SETUP command facility
  show           Show running system information
  sleep          Sleep for the specified number of seconds
  ssh            SSH to another system
  system         System management commands
  tac-pac        Save tac information to a specific location
  tail           Display the last part of a file
  telnet         Telnet to another system
  terminal       Set terminal line parameters
  test           Test command
  traceroute     Trace route to destination
  undebug        Disable Debugging functions (See also debug)
  update         Update license
  write          Write current configuration
  zone           Execute Zone Server commands
  zoneset        Execute zoneset commands
```

Configuration Mode

In configuration mode, you can make changes to the existing configuration. When you save the configuration, these commands are preserved across switch reboots. Once you are in configuration mode, you can enter interface configuration submode, zone configuration submode, and a variety of feature-specific submodes. Configuration mode is the starting point for all configuration commands. When you are in configuration mode, the switch expects configuration commands from the user.

The following example shows output from the **config terminal** command:

```
switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

Configuration Mode Commands and Submodes

Here is a list of configuration mode commands:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ?
Configure commands:
  aaa                   Configure aaa functions
  arp                   [no] remove an entry from the ARP cache
  banner                Configure banner message
  boot                  Configure boot variables
  callhome              Enter the callhome configuration mode
  cdp                   CDP Configuration parameters
  cfs                   CFS configuration commands
  cimserver             Modify cimserver configuration
  cli                   CLI configuration commands
  clock                 Configure time-of-day clock
  cloud-discovery       Configure Cloud Discovery
  crypto                Set crypto settings
  device-alias          Device-alias configuration commands
  do                    EXEC command
  dpvm                  Configure Dynamic Port Vsan Membership
  end                   Exit from configure mode
```

```
exit                  Exit from configure mode
fabric-binding        Fabric Binding configuration
fc-tunnel             Configure fc-tunnel
fcalias               Fcalias configuration commands
fcanalyzer            Configure cisco fabric analyzer
fcc                   Configure FC Congestion Control
fcdomain              Enter the fcdomain configuration mode
fcdroplatency         Configure switch or network latency
fcflow                Configure fcflow
fcid-allocation       Add/remove company id(or OUIs) from auto area list
fcinterop             Interop commands
fcip                  Enable/Disable FCIP
fcns                  Name server configuration
fcroute               Configure FC routes
fcrxbbcredit          Enable extended rx b2b credit configuration
fcs                   Configure Fabric Config Server
fcsp                  Config commands for FC-SP
fctimer               Configure fibre channel timers
fdmi                  Config commands for FDMI
ficon                 Configure ficon information
fspf                  Configure fspf
hw-module             Enable/Disable OBFL log information
in-order-guarantee    Set in-order delivery guarantee
interface             Select an interface to configure
ip                    Configure IP features
ips                   Various sibyte module related commands
ipv6                  Configure IPv6 features
iscsi                 Enable/Disable iSCSI
islb                  ISCSI server load-balancing
isns                  Configure iSNS
isns-server           ISNS server
ivr                   Config commands for IVR
kernel                Kernel options
line                  Configure a terminal line
logging               Modify message logging facilities
mcast                 Configure multicast
nasb                  Configure Third-Party Copy Functionality
no                    Negate a command or set its defaults
ntp                   NTP Configuration
port-security         Configure Port Security
port-track            Configure Switch port track config
power                 Configure power supply
poweroff              Poweroff a module in the switch
qos                   QoS Configuration commands
radius                Configure RADIUS configuration
radius-server         Configure RADIUS related parameters
rib                   Configure RIB parameters
rmon                  Remote Monitoring
role                  Configure roles
rscn                  Config commands for RSCN
san-ext-tuner         Enable/Disable San Extension Tuner tool
santap                Enter SanTap configuration
scheduler             Config commands for scheduler
scsi-target           Scsi-target configuration
snmp-server           Configure snmp server
span                  Enter SPAN configuration mode
ssh                   Configure SSH parameters
ssm                   Config commands for SSM (Storage Services Module)
switchname            Configure system's network name
system                System config command
tacacs+               Enable tacacs+
telnet                Enable telnet
tlport                Configure TL Port information
trunk                 Configure Switch wide trunk protocol
```

```
username              Configure user information.
vsan                  Enter the vsan configuration mode
wwn                   Set secondary base MAC addr and range for additional WWNs
zone                  Zone configuration commands
zone-attribute-group  Zone attribute group commands
zoneset               Zoneset configuration commands
```

Configuration mode, also known as terminal configuration mode, has several submodes. Each of these submodes places you deeper in the prompt hierarchy. When you type **exit**, the switch backs out one level and returns you to the previous level. When you type **end**, the switch backs out to the user EXEC level. You can also type **Ctrl-Z** in configuration mode as an alternative to typing **end**.

Note In configuration mode, you can alternatively enter
— **Ctrl-Z** instead of the **end** command, and
— **Ctrl-G** instead of the **exit** command

You can execute an EXEC mode command from a configuration mode or submode prompt. You can issue this command from any submode within the configuration mode. When in configuration mode (or in any submode), enter the **do** command along with the required EXEC mode command. The entered command is executed at the EXEC level and the prompt resumes its current mode level.

```
switch(config)# do terminal session-timeout 0
switch(config)#
```

In this example, **terminal session-timeout** is an EXEC mode command—you are issuing an EXEC mode command using the configuration mode **do** command.

The **do** command applies to all EXEC mode commands other than the **end** and **exit** commands. You can also use the help (**?**) and command completion (**Tab**) features for EXEC commands when issuing a **do** command along with the EXEC command.

Table 2-2 lists some useful command key combinations that can be used in both EXEC and configuration modes:

Table 2-2 Useful Command Key Combination Descriptions

Command	Description
Ctrl-P	Up history.
Ctrl-N	Down history.
Ctrl-R	Refreshes the current line and reprints it.
Ctrl-X H	List history. When using this key combination, press and release the Ctrl and X keys together before pressing the H key.
Alt-P	History search backwards. **Note** The difference between **Tab** completion and **Alt-P** or **Alt-N** is that **Tab** completes the current word while **Alt- P** and **Alt-N** completes a previously entered command.
Alt-N	History search forwards.
Ctrl-G	Exit.
Ctrl-Z	End.
Ctrl-L	Clear screen.

Displaying the Present Working Context

Some features have configuration submode hierarchy nested more than one level deep. In these cases, you can display the commands you used to reach your present working context (PWC). To display the command used to reach the current PWC, issue the **pwc** command at any command mode prompt.

```
switch(config-if)# pwc
(config t) -> (int mgmt 0)
```

CLI Command Navigation

To redisplay a command you previously entered, press the **Up Arrow** key. You can continue to press the **Up Arrow** key to see more previously issued commands. Similarly, you can press the **Down Arrow**, **Right Arrow**, **Left Arrow**, and **Delete** keys to navigate through the command history and to modify an existing command string.

Command Completion

In any command mode, you can begin a particular command sequence and immediately press the **Tab** key to complete the rest of the command.

```
switch(config)# ro<Tab>
switch(config)# role <Tab>
switch(config)# role name
```

This form of help is called command completion, because it completes a word for you. If several options are available for the typed letters, all options that match those letters are presented:

```
switch(config)# fc<Tab>
fc-tunnel        fcalias          fcanalyzer       fcc
fcdomain         fcdroplatency    fcflow           fcid-allocation
fcinterop        fcip             fcns             fcroute
fcrxbbcredit     fcs              fcsp             fctimer

switch(config)# fcd <Tab>
fcdomain       fcdroplatency
switch(config)# fcdo<Tab>
switch(config)# fcdomain
```

File System Completion

You can use the **Tab** key to complete schemes, servers, and file names available in the file system.

For example,

```
switch# cd bootflash:<Tab>
bootflash:               bootflash://sup-1/        bootflash://sup-remote/
bootflash:///            bootflash://sup-2/        bootflash://sup-standby/
bootflash://module-5/    bootflash://sup-active/
bootflash://module-6/    bootflash://sup-local/

switch# cd bootflash://mo<Tab>
bootflash://module-5/  bootflash://module-6/
cvswitch# cd bootflash://module-
```

The no and Default Forms of Commands

You can issue the **no** form of any command to perform the following actions:

- Undo a wrongly issued command.

  If you issue the **member** command in zone configuration submode, you can undo the results:

```
switch(config)# zone name test vsan 1
switch(config-zone)# member pwwn 12:12:12:12:12:12:12:12
switch(config-zone)# no member pwwn 12:12:12:12:12:12:12:12
WARNING: Zone is empty. Deleting zone test. Exit the submode.
switch(config-zone)#
```

- Delete a created facility.

  If you want to delete a zone that you created:

```
switch(config)# zone name test vsan 1
switch(config-zone)# exit
switch(config)# no zone name test vsan 1
switch(config)#
```

  You cannot delete a zone facility called test while residing in it. You must first exit the zone configuration submode and return to configuration mode.

- Revert to the default value.

  If you issue the **zone merge-control restrict vsan** command, you can undo the results:

```
switch(config)# zone zone merge-control restrict vsan 10
switch(config)# no zone merge-control restrict vsan 10
switch(config)#
```

CLI Command Configuration Options

You can configure the software in one of two ways:

- You can create the configuration for the switch interactively by issuing commands at the CLI prompt.
- You can create an ASCII file containing a switch configuration and then load this file on the required system. You can then use the CLI to edit and activate the file (see the "Working with Configuration Files" section on page 8-1).

Getting Help

In any command mode, you can get a list of available commands by entering a question mark (?).

```
switch# ?
```

To obtain a list of commands that begin with a particular character sequence, type in those characters followed immediately by the question mark (?). Do not include a space.

```
switch# co?
configure copy
```

To list keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help, because it reminds you which keywords or arguments are applicable based on the commands, keywords, and arguments you have already entered.

```
switch# config ?
 terminal Configure the system from the terminal
```

Tip If you are having trouble entering a command, check the system prompt and enter the question mark (?) for a list of available commands. You might be in the wrong command mode or using incorrect syntax.

Managing the Switch Configuration

This section includes the following topics:

Displaying the Switch Configuration

You can view the ASCII form of the configuration file when required. To view the current configuration tree from the EXEC prompt, issue the **show running-config** command. If the running configuration is different from the startup configuration, issue the **show startup-config** command to view the ASCII version of the current startup configuration that was used to boot the switch if a **copy run start** command was not issued after the reboot. Use the **show startup-config** command to view the contents of the current startup configuration.

You can also gather specific information on the entire switch configuration by issuing the relevant **show** commands. Configurations are displayed based on a specified feature, interface, module, or VSAN. Available **show** commands for each feature are briefly described in this section and listed at the end of each chapter.

Examples 2-2 to 2-8 display a few **show** command examples.

Example 2-2 Displays Details on the Specified Interface

```
switch# show interface fc1/1
fc1/1 is up
    Hardware is Fibre Channel, 20:01:ac:16:5e:4a:00:00
    vsan is 1
    Port mode is E
    Speed is 1 Gbps
    Beacon is turned off
    FCID is 0x0b0100
      0 frames input, 0 bytes, 0 discards
      0 runts, 0 jabber, 0 too long, 0 too short
      0 input errors, 0 CRC, 0 invalid transmission words
      0 address id, 0 delimiter
      0 EOF abort, 0 fragmented, 0 unknown class
      0 frames output, 0 bytes, 0 discards
      Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
      Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
```

Example 2-3 Displays the Software and Hardware Version

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2006, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software may be covered under the GNU Public
License or the GNU Lesser General Public License. A copy of
each such license is available at
http://www.gnu.org/licenses/gpl.html and
http://www.gnu.org/licenses/lgpl.html

Software
  BIOS:      version 1.0.8
  loader:    version 1.1(2)
  kickstart: version 2.0(1) [build 2.0(0.6)] [gdb]
  system:    version 2.0(1) [build 2.0(0.6)] [gdb]

  BIOS compile time:       08/07/03
  kickstart image file is: bootflash:///m9500-sf1ek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time:  10/25/2010 12:00:00
  system image file is:    bootflash:///m9500-sf1ek9-mzg.2.0.0.6.bin
  system compile time:     10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:           0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
    Reason: Reset Requested by CLI command reload
    System version: 2.0(0.6)
    Service:
```

Example 2-4 Displays the Running Configuration

```
switch# show running-config
Building Configuration ...
 interface fc1/1
 interface fc1/2
 interface fc1/3
 interface fc1/4
 interface mgmt0
ip address 172.22.95.112 255.255.255.0
no shutdown
vsan database
boot system bootflash:system-237; sup-1
boot kickstart bootflash:boot-237 sup-1
callhome
ip default-gateway 172.22.95.1
switchname switch
trunk protocol enable
username admin password 5 /AFDAMD4B2xK2  role network-admin
```

Note The interface configuration information can be display in multiple entries in the running configuration. See the "Displaying Interface Information" section on page 12-20.

Example 2-5 Displays the Difference Between the Running and Startup Configurations

```
switch# show running-config diff
Building Configuration ...
*** Startup-config
--- Running-config
****************** 1,16 ****
  fcip enable
  ip default-gateway 172.22.91.1
  iscsi authentication none
  iscsi enable
! iscsi import target fc
  iscsi virtual-target name vt
    pWWN 21:00:00:04:cf:4c:52:c1
  all-initiator-permit
--- 1,20 ----
  fcip enable
+ aaa accounting logsize 500
+
+
+
  ip default-gateway 172.22.91.1
  iscsi authentication none
  iscsi enable
! iscsi initiator name junk
  iscsi virtual-target name vt
    pWWN 21:00:00:04:cf:4c:52:c1
  all-initiator-permit
```

Example 2-6 Displays the Configuration for a Specified Interface

```
switch# show running-config interface fc2/9
interface fc2/9
switchport mode E
no shutdown
```

Note The **show running-config interface** command is different from the **show interface** command.

Example 2-7 Displays the Configuration for all Interfaces in a 16-Port Module

```
switch# show running-config interface fc2/10 - 12
interface fc2/10
switchport mode E
no shutdown

interface fc2/11
switchport mode E
no shutdown

interface fc2/12
switchport mode FL
no shutdown
```

Example 2-8 Displays the Configuration Per VSAN

```
switch# show runnning vsan 1
Building Configuration ...
zone name m vsan 1
  member pwwn 21:00:00:20:37:60:42:5c
  member pwwn 21:00:00:20:37:4b:00:a2
zoneset name m vsan 1
  member m
zoneset  activate name m vsan 1
```

Saving a Configuration

Use the **copy running-config startup-config** command to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

See the "Copying Configuration Files" section on page 8-5 and the "Preserving Module Configuration" section on page 11-7.

Clearing a Configuration

Use the **write erase** command to clear a startup configuration. Once this command is issued, the switch's startup configuration reverts to factory defaults. The running configuration is not affected.

Caution The **write erase** command erases the entire startup configuration with the exception of any configuration that affects the loader functionality.

The **write erase boot** command only erases the configuration that affects the loader functionality. The loader functionality configuration includes the boot variables and the mgmt0 IP configuration information (IP address, netmask, and default gateway).

```
switch# write erase boot
This command will erase the boot variables and the ip configuration of interface mgmt 0
```

Displaying Users

Use the **show users** command to display all users currently accessing the switch.

```
switch# show users
admin    pts/7        Jan 12 20:56 (10.77.202.149)
admin    pts/9        Jan 12 23:29 (user1.example.com)
admin    pts/11       Jan 13 01:53 (dhcp-10-10-1-1.example.com)
```

Sending Messages to Users

Use the **send** command to send a message to all active CLI users currently using the switch. This message is restricted to 80 alphanumeric characters with spaces.

This command sends a warning message to all active users about the switch being shut down.

```
switch# send Shutting down the system in 2 minutes. Please log off.

Broadcast Message from admin@excal-112
        (/dev/pts/3) at 16:50 ...
Shutting down the system in 2 minutes. Please log off.
```

Using the ping and ping ipv6 Commands

Use the **ping** command to verify the connectivity of a remote host or server by sending echo messages.

The IPv4 syntax for this command is **ping** *host* or **ping** *ipv4-address*.

```
switch# ping 198.133.219.25
PING 198.133.219.25 (198.133.219.25) 56(84) bytes of data.
64 bytes from 198.133.219.25: icmp_seq=1 ttl=245 time=0.856 ms
64 bytes from 198.133.219.25: icmp_seq=2 ttl=245 time=1.02 ms

--- 198.133.219.25 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.856/0.941/1.027/0.090 ms
```

The IPv6 syntax for this command is **ping ipv6** *host* or **ping ipv6** *ipv6-address*. The following example pings an IPv6 link-local address configured on a specific address.

```
shellfish# ping ipv6 fe80::205:30ff:fe01:a4fa interface gigabitethernet 1/1
PING fe80::205:30ff:fe01:a4fa(fe80::205:30ff:fe01:a4fa) from ::1 gige1-1: 56 data bytes
64 bytes from fe80::205:30ff:fe01:a4fa: icmp_seq=1 ttl=64 time=0.091 ms
64 bytes from fe80::205:30ff:fe01:a4fa: icmp_seq=2 ttl=64 time=0.077 ms
64 bytes from fe80::205:30ff:fe01:a4fa: icmp_seq=3 ttl=64 time=0.080 ms
64 bytes from fe80::205:30ff:fe01:a4fa: icmp_seq=4 ttl=64 time=0.075 ms
64 bytes from fe80::205:30ff:fe01:a4fa: icmp_seq=5 ttl=64 time=0.076 ms
```

To abnormally terminate a ping session, type the **Ctrl-C** escape sequence.

Using the Extended ping and ping ipv6 Commands

The **ping** and **ping ipv6** commands provide additional options to verify the connectivity of a remote host or server. To specify these additional parameters, just type **ping** at the CLI switch prompt and press **Enter**.

Table 2-3 summarizes the syntax and the defaults.

Table 2-3 Options and Defaults for the ping and ping ipv6 Commands

Option	Description	Default
Target IP address	The IPv4 address, IPv6 address, or host name of the destination node to ping.	Not applicable
Repeat count	The number of ping packets to be sent to the destination address.	5 packets
Datagram size	The size of each ping packet in bytes.	100 bytes
Timeout in seconds	The timeout interval before the **ping** or **ping ipv6** command is terminated.	2 seconds

Table 2-3 Options and Defaults for the ping and ping ipv6 Commands (continued)

Option	Description	Default
Extended commands	Specifies if a series of additional commands appear.	No
Sweep range of sizes	The sizes of the echo packets being sent. This option determines the minimum sizes of the MTUs configured on the nodes along the path to the destination address. You can then reduce packet fragmentation performance problems (see the "Configuring the MTU Frame Size" section on page 45-3).	No
Source address or interface	The numeric IP address or the name of the source interface.	Not applicable
Type of service	The quality of service (QoS) in Internet Control Message Protocol (ICMP) datagrams (see the "QoS" section on page 56-3).	0
Set DF bit in IP header	The Path MTU Discovery strategy (see the "Configuring the MTU Frame Size" section on page 45-3).	No
Data pattern	You may specify up to 16 bytes to pad the outgoing packet. This padding is useful when diagnosing data-dependent problems in a network. For example, `ff` fills the outgoing packet with all ones.	0xABCD

The syntax for this command is as follows:

```
switch# ping
Target IP address: 198.133.219.25
Target IP address: 198.133.219.25
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]:
Set DF bit in IP header [n]:
Data pattern [0xABCD]:
Sweep range of sizes [n]:
PATTERN: 0xabcd
PING 198.133.219.25 (198.133.219.25) 100(128) bytes of data.
108 bytes from 198.133.219.25: icmp_seq=1 ttl=245 time=0.600 ms
108 bytes from 198.133.219.25: icmp_seq=2 ttl=245 time=0.614 ms
108 bytes from 198.133.219.25: icmp_seq=3 ttl=245 time=0.872 ms
108 bytes from 198.133.219.25: icmp_seq=4 ttl=245 time=0.558 ms
108 bytes from 198.133.219.25: icmp_seq=5 ttl=245 time=0.570 ms

--- 198.133.219.25 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 7996ms
rtt min/avg/max/mdev = 0.558/0.642/0.872/0.120 ms
```

To abnormally terminate a ping session, type the **Ctrl-C** escape sequence.

Using traceroute and traceroute ipv6 Commands

Use the **traceroute** command to print the routes taken to reach a specified host or IP address.

The IPv4 syntax for this command is **traceroute** *host* or **traceroute** *ipv4-address*.

```
switch# traceroute www.cisco.com
Tracing route to www.cisco.com [198.133.219.25] 30 hops max, 38 byte packets
  1     bras3-l0.pltnca.sbcglobal.net [151.164.184.79] 30 ms    30 ms    20 ms
  2     dist2-vlan50.pltn13.pbi.net [64.164.97.67] 20 ms    20 ms    30 ms
  3     bb2-g1-1.pltn13.pbi.net [67.116.251.194] 20 ms    20 ms    20 ms
  4     bb1-p12-0.pltn13.pbi.net [151.164.40.17] 20 ms    21 ms    20 ms
  5     bb2-p13-0.sntc01.pbi.net [151.164.191.65] 20 ms    20 ms    30 ms
  6     ex1-p3-0.eqsjca.sbcglobal.net [64.161.1.54] 20 ms    20 ms    30 ms
  7     sl-st20-sj-0-0.sprintlink.net [144.223.242.81] 20 ms    20 ms    30 ms
  8     sl-bb25-sj-10-0.sprintlink.net [144.232.20.62] 20 ms    30 ms    20 ms
  9     sl-gw11-sj-10-0.sprintlink.net [144.232.3.134] 70 ms    30 ms    30 ms
 10     sl-ciscopsn2-11-0-0.sprintlink.net [144.228.44.14] 20 ms    30 ms    20 ms
 11     sjce-dmzbb-gw1.cisco.com [128.107.239.89] 20 ms    30 ms    30 ms
 12     sjck-dmzdc-gw1.cisco.com [128.107.224.69] 20 ms    30 ms    20 ms
 13     www.cisco.com (198.133.219.25)  2.496 ms *  2.135 ms
```

The IPv6 syntax for this command is **traceroute ipv6** *host* or **traceroute ipv6** *ipv6-address*.

```
switch# traceroute ipv6
Target IPv6 address: 2001:0DB8::3/64
Datagram size [40]:
Extended commands [n]: y
Maximum time-to-live [30]:
Source address:
Port number [33434]:
```

To cancel a **traceroute** or **traceroute ipv6** command before it completes, enter **Ctrl-C**.

Configuring Terminal Parameters

This section includes the following topics:

Setting the Terminal Session Timeout

Use the **exec-timeout** command in configuration mode to configure the lifetime of all terminal sessions on that switch. When the time limit configured by this command is exceeded, the shell exits and closes that session. The syntax for this command is **exec-timeout** *minutes*.

The default is 30 minutes. You can configure different timeout values for a console or a virtual terminal line (VTY) session. You can set the **exec-timeout** value to 0 to disable this feature so the session remains active until you exit the switch. This change is saved in the configuration file.

- From the console:

```
switch(config)# line console
switch(config-console)# exec-timeout 60
```

  Specifies the current console shell timeout to be 60 minutes.

- From a VTY session (Telnet or SSH):

```
switch(config)# line vty
switch(config-line)# exec-timeout 60
```

  Specifies the current console shell timeout to be 60 minutes.

Displaying Terminal Sessions

Use the **show line** command to display all configured terminal sessions:

```
switch# show line
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Disable
    Modem Init-String -
        default : ATE0Q1&D2&C1S0=1\015
    Statistics:  tx:5558511     rx:5033958     Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Disable
    Modem Init-String -
        default : ATE0Q1&D2&C1S0=1\015
    Hardware Flowcontrol: ON
    Statistics:  tx:35     rx:0     Register Bits:RTS|DTR
```

Clearing Terminal Sessions

Use the **clear line** command to clear a specified terminal session:

```
switch# clear line Aux
```

Setting the Terminal Timeout

Use the **terminal session-timeout** command in EXEC mode to configure the automatic logout time for the current terminal session on that switch. When the time limit configured by this command is exceeded, the switch closes that session and exits.

The syntax for this command is **terminal session-timeout** *minutes*.

The default is 30 minutes. You can set the **terminal session-timeout** value to 0 to disable this feature so the terminal remains active until you choose to exit the switch. This change is not saved in the configuration file.

```
switch# terminal session-timeout 600
```

Specifies the terminal timeout to be 600 minutes for the current session.

Setting the Terminal Type

Use the **terminal terminal-type** command in EXEC mode to specify the terminal type for a switch:

The syntax for this command is **terminal terminal-type** *terminal-type*.

```
switch# terminal terminal-type vt100
```

Specifies the terminal type. The *terminal-type* string is restricted to 80 characters and must be a valid type (for example vt100 or xterm). If a Telnet or SSH session specifies an unknown terminal type, the switch uses the vt100 terminal by default.

Setting the Terminal Screen Length

Use the **terminal length** command in EXEC mode to set the terminal screen length for the current session. This command is specific to only the console port. Telnet and SSH sessions set the length automatically.

The syntax for this command is **terminal length** *lines*.

```
switch# terminal length 20
```

Sets the screen length for the current session to 20 lines for the current terminal session. The default is 24 lines.

Setting the Terminal Screen Width

Use the **terminal width** command in EXEC mode to set the terminal screen width for the current session. This command is specific to only the console port. Telnet and SSH sessions set the width automatically.

The syntax for this command is **terminal width** *columns*.

```
switch# terminal width 86
```

Sets the screen length for the current session to 86 columns for the current terminal session. The default is 80 columns.

Displaying Terminal Settings

Use the **show terminal** command to display the terminal settings for the current session:

```
switch# show terminal
TTY: Type: "vt100"
Length: 24 lines, Width: 80 columns
Session Timeout: 525600 minutes
```

Configuring the Switch Banner Message

You can issue the **banner motd** command in configuration mode to configure a message of the day (MOTD).

The syntax for this command is **banner motd [***delimiting-character message delimiting-character***]**

The following example configures a banner message with the following text "Testing the MOTD Feature."

```
switch# config t
switch(config)# banner motd # Testing the MOTD Feature. #
```

The message is restricted to 40 lines with a maximum of 80 characters in each line.

Use the **show banner motd** command to display the configured banner message.

The following example displays the configured banner message.

```
switch# show banner motd
 Testing the MOTD Feature
```

The configured MOTD banner is displayed before the login prompt on the terminal whenever a user logs in to a Cisco MDS 9000 Family switch.

```
 Testing the MOTD Feature
switch login:
```

Follow these guidelines when choosing your delimiting character:

- Do not use the *delimiting-character* in the *message* string.
- Do not use **"** and **%** as delimiters.

You can include tokens in the form $ (token) in the message text. Tokens will be replaced with the corresponding configuration variable. For example:

- `$(hostname)` displays the host name for the switch
- `$(line)` displays the vty or tty line or name

The following example spans multiple lines and uses tokens to configure the banner message:

```
switch# config t
switch(config)# banner motd #
Enter TEXT message. End with the character '#'.
Welcome to switch $(hostname).
Your tty line is $(line).
#
```

Directing show Command Output to a File

You can direct **show** command output to a file, either on the volatile file system, on slot0 CompactFlash memory, or on a remote server.

The following example shows how to direct the **show running-config** output to a file on the volatile file system.

```
switch1# show running-config > volatile:switch1-run.cfg
```

The following example shows how to direct the **show running-config** output to a file on slot0 CompactFlash memory.

```
switch2# show running-config > slot0:switch2-run.cfg
```

The following example shows how to direct the **show running-config** output to a file on a TFTP server.

```
switch3# show running-config > tftp://10.10.1.1/home/configs/switch3-run.cfg
Preparing to copy...done
```

Using CLI Variables

The SAN-OS CLI parser supports the definition and use of variables in CLI commands. CLI variables can be used as follows:

- Entered directly on the command line.
- Passed to the child script initiated using the **run-script** command. The variables defined in the parent shell are available for use in the child **run-script** command process (see the "Executing Commands Specified in a Script" section on page 2-34).
- Passed as command line arguments to the **run-script** command (see the "Executing Commands Specified in a Script" section on page 2-34).

CLI variables have the following characteristics:

- You cannot reference a variable through another variable using nested references.
- You can define persistent variables that are available across switch reloads.
- You can reference only one predefined system variable, the TIMESTAMP variable.

User-Defined CLI Session Variables

You can define CLI session variables to persist only for the duration of your CLI session using the **cli var name** command in EXEC mode. CLI session variables are useful for scripts that you execute periodically.

The following example shows how to create a user-defined CLI session variable.

```
switch# cli var name testinterface fc 1/1
```

You can reference a variable using the syntax **$**(variable).

The following example shows how to reference a user-defined CLI session variable.

```
switch# show interface $(testinterface)
fc1/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
```

```
    Port WWN is 20:01:00:0d:ec:0e:1d:00
    Admin port mode is auto, trunk mode is on
    snmp traps are enabled
    Port mode is F, FCID is 0x01000b
    Port vsan is 1
    Speed is 2 Gbps
    Transmit B2B Credit is 7
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 256 bits/sec, 32 bytes/sec, 1 frames/sec
    5 minutes output rate 256 bits/sec, 32 bytes/sec, 1 frames/sec
      232692 frames input, 7447280 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      232691 frames output, 7448692 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      7 transmit B2B credit remaining
```

Use the **show cli variables** command to display user-defined CLI session variables.

The following example displays user-defined CLI session variables.

```
switch# show cli variables
VSH Variable List
-----------------
TIMESTAMP="2005-10-24-21.29.33"
testinterface="fc 1/1"
```

Use the **cli no var name** command to remove user-defined CLI session variables.

The following example removes a user-defined CLI session variable.

```
switch# cli no var name testinterface
```

User-Defined CLI Persistent Variables

You can define CLI variables that persist across CLI sessions and switch reloads using the **cli var name** command in configuration mode. These CLI persistent variables are configured in configuration mode and are saved in the running configuration file.

The following example shows how to create a user-defined CLI persistent variable.

```
switch# config t
switch(config)# cli var name mgmtport mgmt 0
switch(config)# exit
switch#
```

You can reference a variable using the syntax **$**(variable).

The following example shows how to reference a user-defined CLI persistent variable.

```
switch# show interface $(mgmtport)
mgmt0 is up
    Hardware is FastEthernet
    Address is 000e.38c6.2c6c
    Internet address is 10.10.10.1/24
    MTU 1500 bytes, BW 100 Mbps full Duplex
    288996 packets input, 97746406 bytes
```

```
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    9089 packets output, 1234786 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Use the **show cli variables** command to display user-defined CLI persistent variables.

The following example displays user-defined CLI persistent variables.

```
switch# show cli variables
VSH Variable List
-----------------
TIMESTAMP="2005-10-24-21.37.13"
mgmtport="mgmt 0"
```

Use the **no cli var name** command in configuration mode to remove user-defined CLI persistent variables.

The following example removes a user-defined CLI persistent variable.

```
switch# config t
switch(config)# no cli var name mgmtport
```

System-Defined Variables

Cisco MDS SAN-OS supports one predefined variable: TIMESTAMP. This variable refers to the time of execution of the command in the format YYYY-MM-DD-HH.MM.SS.

Note The TIMESTAMP variable name is case sensitive. All letters must be uppercase.

The following example uses $(TIMESTAMP) when periodically gathering statistics into files using the command scheduler.

```
switch# config t 1
switch(config)# scheduler enable
switch(config)# scheduler logfile size 16
switch(config)# scheduler job name j1
switch(config-job)# show interface mgmt0 | include mgmt > file
switch(config-job)# copy volatile:file bootflash:file.$(TIMESTAMP)
switch(config-job)# end
switch(config)#
```

The following example uses $(TIMESTAMP) when redirecting **show** command output to a file.

```
switch# show running-config > rcfg.$(TIMESTAMP)
Preparing to copy....done
switch# dir volatile:
       7231     Oct 03 20:20:42 2005  rcfg.2005-10-03-20.20.42

Usage for volatile://sup-local
8192 bytes used
20963328 bytes free
20971520 bytes total
```

Using Command Aliases

Command alias support has the following characteristics:

- Command aliases are global for all user sessions.
- Command aliases are persist across reboots.
- Commands being aliased must be typed in full without abbreviation.
- Command alias translation always takes precedence over any keyword in any configuration mode or submode.
- Command alias support is only available on the supervisor module, not the switching modules.
- Command alias configuration takes effect for other user sessions immediately.
- You cannot override the default command alias **alias,** which aliases the **show cli alias**.
- Nesting of command aliases is permitted to a maximum depth of 1. One command alias can refer to another command alias that must refer to a valid command, not to another command alias.
- A command alias always replaces the first command keyword on the command line.
- You can define command aliases for commands in any configuration submode or the EXEC mode.

Defining Command Aliases

You can define command aliases using the **cli alias name** command in configuration mode.

This following example shows how to define command aliases.

```
switch# config t
switch(config)# cli alias name gigint interface gigabitethernet
switch(config)# cli alias name shintbr show interface brief
switch(config)# cli alias name shfcintup "shintbr| include up | include fc"
```

You can display the command aliases defined on the switch using the **alias** default command alias.

The following example shows how to display the command aliases defined on the switch.

```
switch# alias
CLI alias commands
==================
alias      :show cli alias
gigint     :interface gigabitethernet
shintbr    :show interface brief
shfcintup  :shintbr | include up | include fc
```

About Flash Devices

Every switch in the Cisco MDS 9000 Family contains one internal bootflash (see Figure 2-2). The Cisco MDS 9500 Series additionally contains one external CompactFlash called slot0 (see Figure 2-2 and Figure 2-3).

Figure 2-2 Flash Devices in the Cisco MDS 9000 Supervisor Module

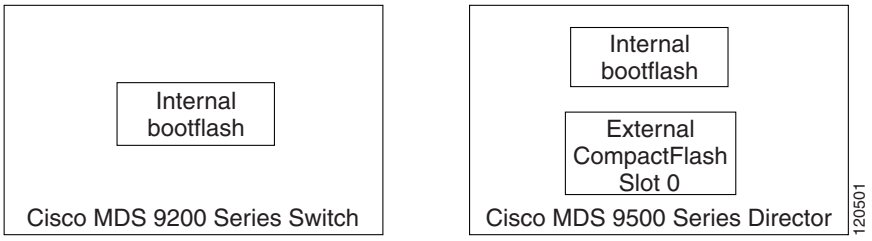

Figure 2-3 External CompactFlash in the Cisco MDS 9000 Supervisor Module

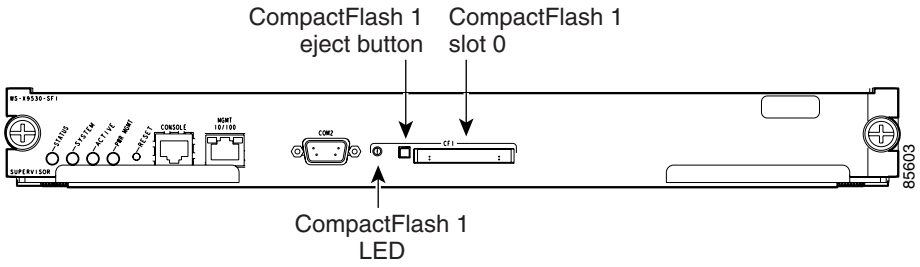

Internal bootflash

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module.You have access to two locations within the internal bootflash: file system.

- The volatile: file system provides temporary storage, and it is also the default location for file system commands. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash: (nonvolatile storage) file system provides permanent storage. The files in bootflash: are preserved through reboots and power outages.

External CompactFlash (Slot0)

Cisco MDS 9500 Series directors contain an additional external CompactFlash referred to as the slot0: file system.

The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

Formatting Flash Devices and File Systems

By formatting a Flash device or a file system, you are clearing out the contents of the device or the file system and restoring it to its factory-shipped state.

See the "About Flash Devices" section on page 2-24 and the "Using Switch File Systems" section on page 2-27.

Initializing Internal bootflash

When a switch is shipped, the **init system** command is already performed and you do not need to issue it again. Initializing the switch resets the entire internal Flash device and erases all data in the bootflash: file system. The internal Flash device is composed of several file systems with bootflash: being one of them. All files in bootflash: are erased and you must download the system and kickstart images again. After issuing an **init system** command, you do not have to format the bootflash: again because bootflash: is automatically formatted.

Note The **init system** command also installs a new loader from the existing (running) kickstart image. You can access this command from the `switch(boot)#` prompt (see Chapter 7, "Software Images").

Caution If your system has an active supervisor module currently running, you must issue the **system standby manual-boot** command in EXEC mode on the active supervisor module before issuing the **init system** command on the standby supervisor module to avoid corrupting the internal bootflash:. After the **init system** command completes on the standby supervisor module, issue the **system no standby manual-boot** command in EXEC mode on the active supervisor module.

If bootflash: is found corrupted during a boot sequence, you will see the following message:

```
ERROR:bootflash: has unrecoverable error; please do "format bootflash:"
```

Use the **format bootflash:** command to only format the bootflash: file system. You can issue the **format bootflash:** command from either the `switch#` or the `switch(boot)#` prompts.

If you issue the **format bootflash:** command, you must download the kickstart and system images again.

Formatting External CompactFlash

Be sure to format an external CompactFlash device before using it to save files or images.

You can verify that the external CompactFlash device is formatted by inserting it into slot0: and issuing the **dir slot0:** command.

- If the external CompactFlash device is already formatted, you can see file system usage information (along with any existing files).
- If the external CompactFlash device is unformatted (corrupted), you will see the following message:

  ```
  Device unavailable
  ```

  In this case, you need to format the CompactFlash device using the **format slot0:** command.

Note The slot0: file system cannot be accessed from either the standby `loader>` prompt or the `switch(boot)#` prompt if the disk is inserted after booting the switch.

Caution The Cisco SAN-OS software only supports CompactFlash devices that are certified by Cisco Systems and formatted using Cisco MDS switches. Using uncertified CompactFlash devices may result in unpredictable consequences; formatting CompactFlash devices using other platforms may result in errors.

Using Switch File Systems

The switch provides the following useful functions to help you manage software image files and configuration files:

Specifying File Systems

The syntax for specifying a file system is *scheme*:[//*server*/]. Table 2-4 describes the file system syntax components.

Table 2-4 File System Syntax Components

Scheme	Server	Description
bootflash	sup-active sup-local sup-1 module-5[1] module-7[2]	Internal CompactFlash memory located on the active supervisor used for storing system images, configuration files, and other miscellaneous files.
	sup-standby sup-remote sup-2 module-6[1] module-8[2]	Internal CompactFlash memory located on the standby supervisor used for storing system images, configuration files, and other miscellaneous files.
slot0	—	External CompactFlash installed in a supervisor module used for storing system images, configuration files, and other miscellaneous files
volatile	—	Volatile random-access memory (VRAM) located on a supervisor module used for temporary or pending changes

Table 2-4 File System Syntax Components (continued)

Scheme	Server	Description
nvram	—	Nonvolatile random-access memory (NVRAM) located on a supervisor module used for storing the startup-config file
log	—	Memory on the active supervisor that stores logging file statistics
system	—	Memory on a supervisor module used for storing the running-config file
modflash	slot-*slot*	CompactFlash located on a Storage Services Module (SSM) used for storing the SSI boot image

1. Cisco MDS 9506 and Cisco MDS 9509 switches
2. Cisco MDS 9513 Directors

Setting the Current Directory

The **cd** command changes the current directory level to a specified directory level. CLI defaults to the volatile: file system. This command expects a directory name input.

Tip Any file saved in the volatile: file system is erased when the switch reboots.

The syntax for this command is **cd** *directory name*

This example changes the current directory to the root directory on the bootflash: file system.

```
switch# cd bootflash:
```

This example changes the current directory to the mydir directory that resides in the slot0: file system.

```
switch# cd slot0:mydir
```

This example changes the current directory to the mystorage directory that resides in the current directory.

```
switch# cd mystorage
```

If the current directory is slot0:mydir, this command changes the current directory to slot0:mydir/mystorage.

Displaying the Current Directory

The **pwd** command displays the current directory location. This example changes the directory and displays the current directory.

```
switch# cd bootflash:
switch# pwd
bootflash:
```

Note If you issue this command from the active supervisor module in a Cisco MDS 9500 Series (for example, `module-5`), then you cannot change the current working directory to the bootflash: of `module-6`. See the "Supervisor Modules" section on page 11-2.

Displaying File Checksums

The **show file** *file* **md5sum** command provides the MD5 checksum of the file. MD5 is an electronic fingerprint for the file. MD5 is the latest implementation of the Internet standards described in RFC 1321 and is useful for data security as well as integrity.

The **show file** *file* **cksum** command provides the checksum of the file. The checksum values compute a cyclic redundancy check (CRC) for each named file. Use this command to verify that the files are not corrupted—compare the checksum output for the received file against the checksum output for the original file.

This example provides the output of the **show file** command when a file is specified.

```
switch# show file bootflash://sup-1/ultimate_file.tar cksum
2569913991

switch# show file bootflash://sup-1/ultimate_file.tar md5sum
52479aae2dce1fd849b6f4916d750392
```

Listing the Files in a Directory

The **dir** command displays the contents of the current directory or the specified directory. The syntax for this command is **dir** *directory* or **dir** *filename*.

This example shows how to list the files on the default volatile: file system.

```
switch# dir
           Usage for volatile: filesystem
                          0 bytes total used
                   20971520 bytes free
                   20971520 bytes available
```

Creating a Directory

The **mkdir** command creates a directory at the current directory level or at a specified directory level.

The syntax for this command is **mkdir** *directory name.*

This example creates a directory called test in the slot0 directory.

```
switch# mkdir slot0:test
```

This example creates a directory called test at the current directory level.

```
switch# mkdir test
```

If the current directory is slot0:mydir, this command creates a directory called slot0:mydir/test.

Deleting an Existing Directory

The **rmdir** command deletes an existing directory at the current directory level or at a specified directory level. The directory must be empty to be deleted.

The syntax for this command is **rmdir** *directory name.*

This example deletes the directory called test in the slot0 directory.

```
switch# rmdir slot0:test
This is a directory.  Do you want to continue (y/n)?  [y] y
```

The **delete** command is also capable of deleting empty and non-empty directories. When you issue this command a warning is displayed to confirm your intention to delete the directory.

This example deletes the directory called test at the current directory level.

```
switch# rmdir test
This is a directory.  Do you want to continue (y/n)?  [y] y
```

If the current directory is slot0:mydir, this command deletes the slot0:mydir/test directory.

Moving Files

The **move** command removes a file from the source directory and places it in the destination directory.

Caution If a file with the same name already exists in the destination directory, that file is overwritten by the moved file.

This example moves the file called samplefile from the root directory of the slot0: file system to the mystorage directory.

```
switch# move slot0:samplefile slot0:mystorage/samplefile
```

This example moves a file from the current directory level.

```
switch# move samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command moves slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

Copying Files

The **copy** command copies a file between file systems within a switch.

Note Use the **dir** command to ensure that enough space is available in the target file system. If enough space is not available, use the **delete** command to remove unneeded files.

The syntax for the **copy** command follows and is explained in Table 2-5.

```
switch# copy scheme:[//module/]filename scheme:[//module/]filename
```

Table 2-5 copy Command Syntax

Scheme	Module	File Name
bootflash	sup-active sup-standby sup-1, module-5[1], or module-7[2] sup-2, module-6[1], or module-8[2] sup-local sup-remote	User-specified
slot0	—	User-specified
volatile	—	User-specified
nvram	—	startup-config or snapshot-config
system	—	running-config

1. Cisco MDS 9506 and Cisco MDS 9509 switches
2. Cisco MDS 9513 Directors

This example copies the file called samplefile from the root directory of the slot0: file system to the mystorage directory.

```
switch# copy slot0:samplefile slot0:mystorage/samplefile
```

This example copies a file from the current directory level.

```
switch# copy samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command copies slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

This example shows how to copy a file from the active supervisor module's (sup-1 in slot 5 on the Cisco MDS 9506 and Cisco MDS 9509 switches or slot 7 on the Cisco MDS 9513 switch) bootflash to the standby supervisor module's (sup-2 in slot 6 on the Cisco MDS 9506 and Cisco MDS 9509 switches or slot 7 on the Cisco MDS 9513 switch) bootflash.

```
switch# copy bootflash:system_image bootflash://sup-2/system_image
```

This example shows how to overwrite the contents of an existing configuration in NVRAM.

```
switch# copy nvram:snapshot-config nvram:startup-config
Warning: this command is going to overwrite your current startup-config.
Do you wish to continue? {y/n} [y] y
```

You can also use the **copy** command to upload and download files from the slot0: or bootflash: file system to or from a FTP, TFTP, SFTP, or SCP server (see the "Copying Configuration Files" section on page 8-5).

Deleting Files

The **delete** command deletes a specified file or the specified directory and all its contents (see the "Deleting Configuration Files" section on page 8-8).

This example shows how to delete a file from the current working directory.

```
switch# delete dns_config.cfg
```

This example shows how to delete a file from an external CompactFlash (slot0).

```
switch# delete slot0:dns_config.cfg
```

This example deletes the entire `my-dir` directory and all its contents.

```
switch# delete bootflash:my-dir
```

Caution If you specify a directory, the **delete** command deletes the entire directory and all its contents.

Displaying File Contents

The **show file** command displays the contents of a specified file in the file system.

The syntax for this command is **show file** *filename*.

This example displays the contents of the test file that resides in the slot0 directory.

```
switch# show file slot0:test
config t
Int fc1/1
no shut
end
show int fc1/1
```

This example displays the contents of a file residing in the current directory.

```
switch# show file myfile
```

Saving Command Output to a File

You can force all screen output to go to a file by appending **>** *filename* to any command. For example, enter **show interface > samplefile** at the EXEC mode switch prompt to save the interface configuration to *samplefile*—a file created at the same directory level. At the EXEC mode switch prompt, issue a **dir** command to view all files in this directory, including the recently saved *samplefile*. See Chapter 5, "Initial Configuration," for information on saving and copying configuration files, and Chapter 7, "Software Images," for information on saving and copying software images.

Note Redirection is allowed only if the current directory is on the `volatile:` (default) or `slot0:` file systems. Redirection is not allowed if the current directory is on the `bootflash:` file system. The current directory can be viewed using the **pwd** command and changed using the **cd** command.

Compressing and Uncompressing Files

The **gzip** command compresses (zips) the specified file using LZ77 coding.

This example directs the output of the **show tech-support** command to a file (Samplefile) and then zips the file and displays the difference in the space used up in the volatile: directory.

```
switch# show tech-support > Samplefile
Building Configuration ...
switch# dir
    1525859     Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
   19443712 bytes free
   20971520 bytes total
switch# gzip volatile:Samplefile
switch# dir
     266069     Jul 04 00:51:03 2003 Samplefile.gz
Usage for volatile://
     266240 bytes used
   20705280 bytes free
   20971520 bytes total
```

The **gunzip** command uncompresses (unzips) LZ77 coded files.

This example unzips the file that was compressed in the previous example.

```
switch# gunzip samplefile
switch# dir
    1525859     Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
   19443712 bytes free
   20971520 bytes total
```

Displaying the Last Lines in a File

The **tail** command displays the last lines (tail end) of a specified file.

The syntax for this command is **tail** *filename* [*number-of-lines*].

```
switch# tail mylog 10
```

You see the last 10 lines of the mylog file.

Command Scripts

This section includes the following sections:

Executing Commands Specified in a Script

The **run-script** command executes the commands specified in a file. To use this command, be sure to create the file and specify commands in the required order.

Note You cannot create the script files at the switch prompt. You can create the script file on an external machine and copy it the bootflash: directory. This section assumes that the script file resides in the bootflash: directory.

The syntax for this command is **run-script** *filename*.

This example displays the CLI commands specified in the testfile that resides in the slot0 directory.

```
switch# show file slot0:testfile
conf t
interface fc 1/1
no shutdown
end
sh interface fc1/1
```

This file output is in response to the **run-script** command executing the contents in the testfile file:

```
switch# run-script slot0:testfile
'conf t'
Enter configuration commands, one per line. End with CNTL/Z.
'interface fc1/1'
'no shutdown'
'end'
'sh interface fc1/1'
fc1/1 is down (Fcot not present)
    Hardware is Fibre Channel
    Port WWN is 20:01:00:05:30:00:48:9e
    Admin port mode is auto, trunk mode is on
    vsan is 1
    Beacon is turned off
    Counter Values (current):
      0 frames input, 0 bytes, 0 discards
      0 runts, 0 jabber, 0 too long, 0 too short
      0 input errors, 0 CRC, 0 invalid transmission words
      0 address id, 0 delimiter
      0 EOF abort, 0 fragmented, 0 unknown class
      0 frames output, 0 bytes, 0 discards
      Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
      Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
    Counter Values (5 minute averages):
...
```

Using CLI Variables in Scripts

You can use CLI variables defined by the **cli var** command (see the "Using CLI Variables" section on page 2-21) or passed as arguments in the **run-script** command.

The following example shows how to use CLI session variables in a script file used by the **run-script** command.

```
switch# cli var name testinterface fc 1/1

switch# show file bootflash:test1.vsh
show interface $(testvar)

switch# run-script bootflash:test1.vsh
`show interface $(testvar)`
fc1/1 is down (SFP not present)
Hardware is Fibre Channel
Port WWN is 20:01:00:05:30:00:8e:1e
Admin port mode is auto, trunk mode is on
Port vsan is 1
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 frames input, 128 bytes
0 discards, 0 errors
0 CRC, 0 unknown class
0 too long, 0 too short
1 frames output, 128 bytes
0 discards, 0 errors
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
0 output OLS, 0 LRR, 0 NOS, 0 loop inits
0 receive B2B credit remaining
0 transmit B2B credit remaining
```

The following example shows how you can pass CLI session variable as arguments to a child **run-script** command process.

```
switch# show file bootflash:test1.vsh
show interface $(var1) $(var2)

switch# run bootflash:test2.vsh var1="fc1/1" var2="brief"
`show interface $(var1) $(var2)`
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                          Mode  Speed  Channel
                         Mode                                 (Gbps)
--------------------------------------------------------------------------------
fc1/1      1      auto   on      sfpAbsent       --     --           --
```

Setting the Delay Time

The **sleep** command delays an action by a specified number of seconds.

The syntax for this command is **sleep** *seconds*.

```
switch# sleep 30
```

You will see the switch prompt return after 30 seconds.

This command is useful within scripts. For example, if you create a command script called test-script.

```
switch# show file slot0:test-script
discover scsi-target remote
sleep 10
show scsi-target disk
switch# run-script slot0:test-script
```

When you execute the slot0:test-script command script, the switch software executes the **discover scsi-target remote** command, and then waits for 10 seconds before executing the **show scsi-target disk** command.

Send documentation comments to mdsfeedback-doc@cisco.com

PART 2

Cisco MDS SAN-OS Installation and Switch Management

CHAPTER 3

Obtaining and Installing Licenses

Licenses are available in all switches in the Cisco MDS 9000 Family. Licensing allows you to access specified premium features on the switch after you install the appropriate license for that feature.

This chapter contains information related to licensing types, options, procedures, installation, and management for the Cisco MDS SAN-OS software.

This chapter includes the following sections:

Licensing Terminology

The following terms are used in this chapter:

- Licensed feature—Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- Licensed application—A software feature that requires a license to be used.

- License enforcement—A mechanism that prevents a feature from being used without first obtaining a license.
- Node-locked license—A license that can only be used on a particular switch using the switch's unique host ID.
- Host IDs—A unique chassis serial number that is specific to each Cisco MDS switch.
- Proof of purchase—A document entitling its rightful owner to use licensed feature(s) on one Cisco MDS switch as described in that document. Also known as the claim certificate.
- Product Authorization Key (PAK)—The PAK allows you to obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through e-mail.
- License key file—A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
- Counted license—The number of licenses issued for a single feature (for example, FCIP). You can increase counted licenses (incremental licenses) should a need arise in the future.
- Missing license—If the bootflash has been corrupted or a supervisor module replaced after a license has been installed, that license will show as "missing." The feature will still work, but the license count will be inaccurate. You should reinstall the license as soon as possible.
- Incremental license—An additional licensed feature that was not in the initial license file. License keys are incremental—if you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- Port Activation license—A license that activates additional ports on any of the following:
  - Cisco MDS 9124 Multilayer Fabric Switch
  - Cisco MDS 9134 Multilayer Fabric Switch
  - Cisco Fabric Switch for HP c-Class BladeSystem
  - Cisco Fabric Switch for IBM BladeCenter

  For more information refer to Chapter 4, "On-Demand Port Activation Licensing."
- Evaluation license—A temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- Permanent license—A license that is not time bound is called a permanent license.
- Grace period—The amount of time the features in a license package can continue functioning without a license.
- Support—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Licensing Model

Any feature not included in a license package is bundled with the Cisco MDS 9000 Family switches and is provided at no extra charge.

We recommend that you do not download more licenses than can be used for a module or switch.

See Chapter 4, "On-Demand Port Activation Licensing" for information about on-demand port activation licensing.

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licenses allow features that are applicable to the entire switch. The cost varies based on a per-switch usage. Table 3-1 lists the feature-based license packages.
- Module-based licenses allow features that require additional hardware modules. The cost varies based on a per-module usage. An example is the IPS-8 or IPS-4 module using the FCIP feature.

Note Each module requires its own separate license. If you replace a module that requires a license with a module of the same type (such as replacing a Storage Services Module (SSM) with another SSM), the existing license will support the new module.

Note The Cisco MDS 9216i switch enables SAN extension features on the two fixed IP services ports only. The features enabled on these ports are identical to the features enabled by the SAN extension over IP license on the 14/2-port Multiprotocol Services (MPS-14/2) module. If you install a module with IP ports in the empty slot on the Cisco MDS 9216i, a separate SAN extension over IP license is required to enable related features, such as FCIP, on the IP ports of the additional module.

Table 3-1 Feature-Based Licenses

Feature License	Features
Enterprise package (ENTERPRISE_PKG)	• Enhanced security features: – LUN zoning – Read-only zones • FC Port security • VSAN-based access control • Fibre Channel Security Protocol (FC-SP) authentication • Advanced traffic engineering—quality of service (QoS) • IP security (IPsec) protocol for iSCSI and FCIP using the MPS-14/2 module or Cisco MDS 9216i switch • IPsec & IKE for IPv4 • IKE digital certificates • Extended credits using the MPS-14/2 module or the Cisco MDS 9216i Switch • Enhanced VSAN routing—inter-VSAN routing (IVR) over Fibre Channel • IVR Network Address Translation (NAT) over Fibre Channel • Zone-based traffic prioritizing • Zone-based FC QoS • Extended BB_Credits • Fibre Channel write acceleration • SCSI flow statistics • FCIP encryption • Fabric binding for Fibre Channel • SAN device virtualization

Table 3-1 Feature-Based Licenses (continued)

Feature License	Features
SAN extension over IP package for IPS-8 modules (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules (SAN_EXTN_OVER_IP_IPS4)	The following features apply to IPS-8 and IPS-4 modules: • FCIP • FCIP compression • FCIP write acceleration • FCIP tape read acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP • Network Stimulator
SAN extension over IP package for MPS-14/2 modules (SAN_EXTN_OVER_IP_IPS2) **Note** The FCIP, IVR, and SAN extension tuner features are bundled with the Cisco MDS 9216i switch and do not require the SAN extension over IP package to be installed for the fixed IP ports on the integrated supervisor module. You must install a SAN extension over IP package if you install an MPS-14/2, IPS-8, or IPS-4 module in the Cisco MDS9216i switch.	The following features apply to the MPS-14/2 module and the fixed Cisco MDS 9216i IP ports: • FCIP • Hardware-based FCIP compression • FCIP write acceleration • FCIP tape read acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP
SAN extension over IP package for one MPS-18/4 or one MPS-18/4 FIPS in the Cisco MDS 9500 series (SAN_EXTN_OVER_IP_18_4)	The following features apply to the MPS-18/4 or MPS-18/4 FIPS modules: • FCIP • Hardware-based FCIP compression • FCIP write acceleration • FCIP tape read acceleration • SAN extension tuner features • IVR over FCIP • IVR NAT over FCIP

Table 3-1 Feature-Based Licenses (continued)

Feature License	Features
Mainframe package (MAINFRAME_PKG)	• FICON protocol and CUP management • FICON VSAN and intermixing • Switch cascading • Fabric binding for FICON • IBM TotalStorage Virtual Tape Server (VTS) • IBM TotalStorage XRC application • FICON tape acceleration • FICON license for 9100 • Persistent FCIDs for FICON • Config locking for FICON • Port swap, block, prohibit • FICON Qualification
Fabric Manager Server package (FM_SERVER_PKG)	• Centralized, Multiple physical fabric management • Fabric discovery services • Continuous MDS health and event monitoring • Long term historical Fibre Channel performance monitoring and reporting • Custom performance reports and charting for hotspot analysis • Historical Performance Monitoring • Performance prediction • Performance threshold monitoring • Fabric Manager Web Client for operational view • Fabric Manager server proxy services • Server performance summary report • Configurable RRD collection parameters • Data collection auto update • Event forwarding • Filtering by user-defined groups • Custom Reports Enhancements • Fabric Analysis Report • Threshold Configuration Flexibility • Web-based operational view • Roaming User Profiles • Traffic Analyzer for SCSI Flow Statistics

Table 3-1 Feature-Based Licenses (continued)

Feature License	Features
Storage Services Enabler package (STORAGE_SERVICES_ENABLER_PKG)	• The underlying infrastructure and programmatic interface to enable network-hosted storage applications when used with the Storage Services Modules (SSMs). • The intelligent fabric applications running on the SSM that require the SSE license are as follows: – SANTap – Network-Accelerated Serverless Backup (NASB) – Network-based Storage Virtualization – Third-party partner application
On-demand Port Activation Licensing package (PORT_ACTIVATION_PKG) **Note** License Manager does not prevent installing more port licenses then the available physical ports on the switch. The extra licenses if installed, will not affect the normal behaviour of the licensed ports.	• Activates ports (in 8-port increments) on the Cisco MDS 9124 Fabric Switch, which has 24 ports. The first 8 ports are licensed by default. • Activates 8 ports of 4Gbps on the Cisco MDS 9134 Fabric Switch. The switch has 32 ports, 24 of which are licensed by default. • On the Cisco Fabric Switch for HP c-Class BladeSystem, any eight internal ports and external ports ext1 through ext4 are licensed by default. • On the Cisco Fabric Switch for IBM BladeCenter, any seven internal ports and external ports ext0, ext15 and ext16 are licensed by default. See Chapter 4, "On-Demand Port Activation Licensing" for information about on-demand port activation licensing.
10 Gbps Port Activation Package 10G_PORT_ACTIVATION_PKG	• Activates the two 10 Gbps ports on the Cisco MDS 9134 Multilayer Fabric Switch.
Storage Media Encryption (SME) • SME_FOR_IPS_184_PKG • SME_FOR_9222i_PKG	• Activates Storage Media Encryption for Intrusion Prevention System (IPS) Sensor of 184 unit specification. • Activates Storage Media Encryption for MDS 9222i switch.

Table 3-1 Feature-Based Licenses (continued)

Feature License	Features
Data Mobility Manager (DMM) • DMM_FOR_SSM_PKG	• The Cisco MDS 9000 DMM feature runs on the Storage Service Module (SSM) in a MDS series switch. This license will activate Data Mobility Manager (DMM) for Storage Service Module. • Online migration of heterogenous arrays • Simultaneous migration of multiple LUNs • Unequal size LUN migration • Rate adjusted migration • Verification of migrated data • Secure erasure of migrated data • Dual fabric support

Note License packages for Cisco DMM (Cisco Data Mobility Manager) and Cisco SME (Cisco Storage Media Encryption) are documented in the *Cisco MDS Data Mobility Manager Configuration Guide*, and the *Cisco Storage Media Encryption Configuration Guide*.

Licensing High Availability

As with other Cisco MDS SAN-OS features, the licensing feature also maintains the following high availability standards for all switches in the Cisco MDS 9000 Family:

- Installing any license in any switch is a nondisruptive process.
- Installing a license automatically saves a copy of permanent licenses to the chassis in all switches.
- Enabling a license feature without a license key starts a counter on the grace period. You then have 120 days to install the appropriate license keys or disable the use of that feature. If at the end of the 120-day grace period the switch does not have a valid license key for the feature, the feature is automatically disabled by the switch.

Directors in the Cisco MDS 9500 Series have the following additional high availability features:

- The license software runs on both supervisor modules and provides failover protection.
- The license key file is mirrored on both supervisor modules. Even if both supervisor modules fail, the license file continues to function from the version that is available on the chassis.

Options to Install a License

If you have purchased a new switch through either your reseller or through Cisco Systems, you can:

- Obtain a factory-installed license (only applies to new switch orders).
- Perform a manual license installation (applies to existing switches).

Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new switch.

To obtain a factory-installed license for a new Cisco MDS switch, follow these steps:

Step 1 Contact your reseller or Cisco representative and request this service.

Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.

Step 2 Obtain the host ID from the proof of purchase document for future use.

Step 3 Start to use the switch and the licensed features.

Performing a Manual Installation

If you have existing switches or if you wish to install the licenses on your own, you must first obtain the license key file and then install that file in the switch (see Figure 3-1).

Figure 3-1 Obtaining a License Key File

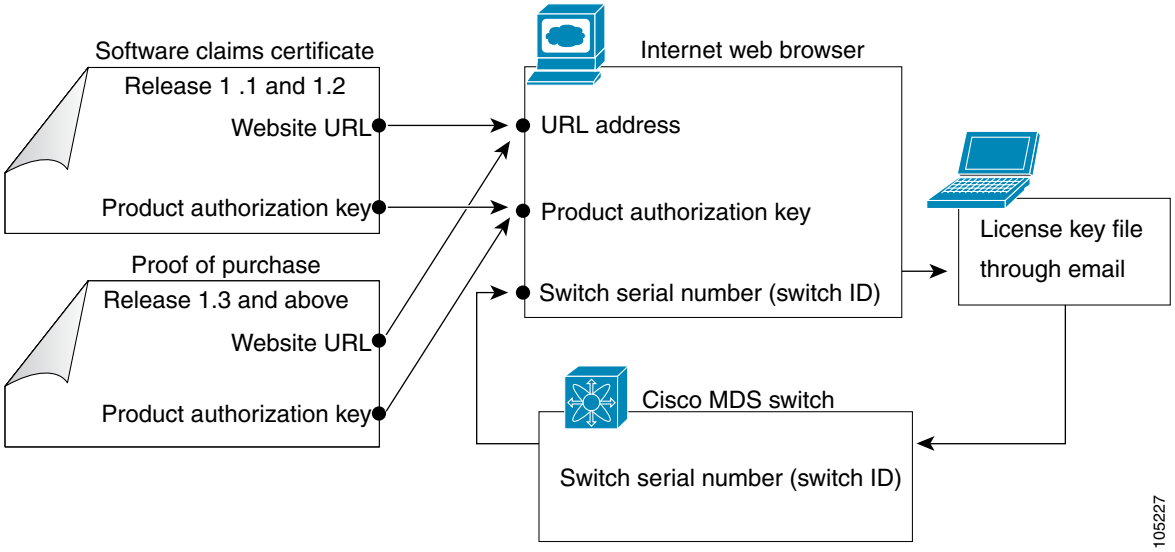

Obtaining the License Key File

Note Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for details on installing automated licenses using the Fabric Manager GUI.

To obtain new or updated license key files using the CLI, follow these steps:

Step 1 Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.

```
switch# show license host-id
License hostid: FOX064317SQ
```

Tip Use the entire ID that appears after the colon (:) sign. In this example, the host ID is `FOX064317SQ`.

Step 2 Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco MDS switch.

Step 3 Obtain the Product Authorization Key (PAK) from either the claim certificate or the proof of purchase document.

Step 4 Locate the website URL from either the claim certificate or the proof of purchase document.

Step 5 Access the specified URL that applies to your switch and enter the switch serial number and the PAK.

The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco SAN-OS software on the specified switch accesses the license key file.

Caution Install the license key file in the specified MDS switch without making any modifications.

A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that feature starts from the first time you start using a feature offered by that license (see the "Grace Period Alerts" section on page 3-15).

Step 6 Use the **copy licenses** CLI command in EXEC mode to save your license file to one of two locations—the bootflash: directory or the slot0: device (see the "Backing Up License Files" section on page 3-12).

Installing the License Key File

Tip If you need to install multiple licenses in any switch in the Cisco MDS 9000 Family, be sure to provide unique file names for each license key file.

Step 7 Select the switches for which you have PAKs or license key files.

When you check the check box for a switch, the PAK or license file name field for that switch becomes editable. The <*serial number*> for each switch is shown in the Host ID column.

Step 8 Enter the PAK or license file name for each switch you have selected in the appropriate column. If you have the license files on your PC, you can double-click in the License File Name text area to bring up a dialog box and browse for the license files.

You can install multiple licenses on the same switch using different PAKs. To do this, enter the PAKs separated by commas.

Step 9 Click **Finish** to transfer the licenses from the host to the switches.

To install a license key file in any switch, follow these steps:

Step 1 Log into the switch through the console port of the active supervisor.

Step 2 Perform the installation by issuing the **install license** command on the active supervisor module from the switch console.

```
switch# install license bootflash:license_file.lic
Installing license ..done
```

Note If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the filename specified in the license key file is used to install the license.

Step 3 Back up the license file to a .tar file on bootflash: using the **copy licenses** command.

```
switch# copy licenses bootflash:/Enterprise.tar
Backing up license done
```

Step 4 Exit the switch console and open a new terminal session to view all license files installed on the switch using the **show license** command.

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
        HOSTID=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Note If the license meets all guidelines when the **install license** command is issued, all features and modules continue functioning as configured. This is true for any switch in the Cisco MDS 9000 Family.

You can use the **show license brief** command to display a list of license files installed on the switch.

```
switch# show license brief
Enterprise.lic
Ficon.lic
FCIP.lic
```

You can use the **show license** *license-name* command to display information about a specific license file installed on the switch.

```
switch# show license file Permanent.lic
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
        HOSTID=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Installing the License Key File to a Remote Location

You can also download the license file to remote locations using the TFTP, SFTP, FTP, or SCP protocols.

Caution Specify the complete path of the remote location. The system will not allow you to proceed if the entire path is not accurately specified. Here are examples of incomplete **install all** commands.

```
switch# install license system bootflash:system-image kickstart tftp
Please provide a complete URI
switch# install license system scp:
Please provide a complete URI
```

Example 3-1 A Sample of the install license Command Issued Using a Remote Download

```
switch# install license bootflash:license_file.lic kickstart tftp:
```

Backing Up License Files

All installed license files can be backed up as a .tar file in the user specified location. Use the **copy licenses** command in EXEC mode to save your license file to one of two locations—bootflash: or slot0:. The following example saves all licenses to a file named Enterprise.tar:

```
switch# copy licenses bootflash:/Enterprise.tar
Backing up license done
```

Tip We recommend backing up your license files immediately after installing them and just before issuing a **write erase** command.

Caution If you erase any existing licenses, you can only install them using the **install license** command.

Identifying License Features in Use

When a Cisco MDS SAN-OS software feature is enabled, it can activate a license grace period. To identify the features active for a specific license, use the **show license usage** *license-name* command.

```
switch# show license usage ENTERPRISE_PKG
```

```
Application
-----------
ivr
qos_manager
-----------
```

Use the **show license usage** command to identify all of the active features on your switch.

```
switch# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 No    -   Unused             Grace 79D 16H
MAINFRAME_PKG                 No    -   Unused             Grace expired
ENTERPRISE_PKG                Yes   -   Unused never       license missing
DMM_FOR_SSM_PKG               No    0   Unused             -
SAN_EXTN_OVER_IP              Yes  16   Unused never       -
PORT_ACTIVATION_PKG           No    0   Unused             -
SME_FOR_IPS_184_PKG           No    0   Unused             Grace 86D 5H
SAN_EXTN_OVER_IP_18_4         No    0   Unused             -
SAN_EXTN_OVER_IP_IPS2         Yes   1   Unused never       1 license(s) missing
SAN_EXTN_OVER_IP_IPS4         No    0   Unused             -
10G_PORT_ACTIVATION_PKG       No    0   Unused             -
SAN_EXTN_OVER_IP_18_4        No    0   Unused             -
STORAGE_SERVICES_ENABLER_PKG  Yes   1   Unused never       1 license(s) missing
--------------------------------------------------------------------------------
```

Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request and issues an error message. Uninstalling an unused license causes the grace period to come into effect. The grace period is counted from the first use of the feature without a license and is reset when a valid license file is installed.

Note Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.

Tip If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.

Caution Disable related features before uninstalling a license. The delete procedure fails if the license is in use.

To uninstall a license, follow these steps:

Step 1 Save your running configuration to a remote server using the **copy** command (see the "Initial Configuration" section on page 5-1).

Step 2 Issue the **show license brief** command in EXEC mode to view a list of all installed license key files and identify the file to be uninstalled. In this example, the file to be uninstalled is the Ficon.lic file.

```
switch# show license brief
Enterprise.lic
Ficon.lic
```

Step 3 Disable the features provided by the license to be uninstalled. Issue the **show license usage** *package_name* command to view the enabled features for a specified package.

```
switch# show license usage ENTERPRISE_PKG
Application
-----------
ivr
qos_manager
-----------
```

Step 4 Uninstall the Enterprise.lic file using the **clear license** *filename* command, where *filename* is the name of the installed license key file.

```
switch# clear license Enterprise.lic
Clearing license Enterprise.lic:
SERVER this_host ANY
VENDOR cisco
```

Step 5 Enter **yes** (yes is the default) to continue with the license update.

```
Do you want to continue? (y/n) y
Clearing license ..done
```

The Enterprise.lic license key file is now uninstalled.

Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.

Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

To update a license, follow these steps:

Step 1 Obtain the updated license file using the procedure described in the "Obtaining the License Key File" section on page 3-10.

Step 2 Save your running configuration to a remote server using the **copy** command (see the "Copying Configuration Files" section on page 8-5).

Step 3 Verify the name of the file to be updated.

Step 4 Follow the procedure for updating a license described in the "Uninstalling Licenses" section on page 3-13.

Step 5 Issue the **show license brief** command to verify the name of the file to be updated.

```
switch# show license brief
sanextn1.lic:
```

Step 6 Update the license file using the **update license** *url* command, where *url* specifies the bootflash:, slot0:, or volatile: location of the updated license file.

```
switch# update license bootflash:sanextn2.lic sanextn1.lic
Updating sanextn1.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=ABCD \
        NOTICE=<LicFileID>san_extn1.lic</LicFileID><LicLineID>0</LicLineID> \
        SIGN=33088E76F668

with bootflash:/sanextn2.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=ABCD \
        NOTICE=<LicFileID>san_extn2.lic</LicFileID><LicLineID>1</LicLineID> \
        SIGN=67CB2A8CCAC2
```

Step 7 Enter **yes** (yes is the default), to continue with the license update.

```
Do you want to continue? (y/n) y
Updating license ..done
switch#
```

The sanextn1.lic license key file is now updated.

Grace Period Alerts

Cisco SAN-OS gives you a 120 day grace period. This grace period starts or continues when you are evaluating a feature for which you have not installed a license.

Note There is no grace period for licenses purchased for the On-Demand Port Activation license feature.

The grace period stops if you disable a feature you are evaluating, but if you enable that feature again without a valid license, the grace period countdown continues where it left off.

The grace period operates across all features in a license package. License packages can contain several features. If you disable a feature during the grace period and there are other features in that license package that are still enabled, the countdown does not stop for that license package. To suspend the grace period countdown for a license package, you must disable every feature in that license package. Use the **show license usage** *license-name* command to determine which applications to disable.

```
switch# show license usage MAINFRAME_PKG
Application
-----------
Ficon
```

```
-----------
```

The Cisco SAN-OS license counter keeps track of all licenses on a switch. If you are evaluating a f feature and the grace period has started, you will receive console messages, SNMP traps, system messages, and Call Home messages on a daily basis.

Beyond that, the frequency of these messages become hourly during the last seven days of the grace period. The following example uses the FICON feature. On January 30th, you enabled the FICON feature, using the 120 day grace period. You will receive grace period ending messages as:

- Daily alerts from January 30th to May 21st.
- Hourly alerts from May 22nd to May 30th.

On May 31st, the grace period ends, and the FICON feature is automatically disabled. You will not be allowed to use FICON until you purchase a valid license.

Note You cannot modify the frequency of the grace period messages.

Caution After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. Any future upgrade will enforce license requirements and the 120-day grace period.

Use the **show license usage** command to display grace period information for a switch.

```
switch# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 No    -   Unused             Grace 79D 16H
MAINFRAME_PKG                 No    -   Unused             Grace expired
ENTERPRISE_PKG                Yes   -   Unused never       license missing
DMM_FOR_SSM_PKG               No    0   Unused             -
SAN_EXTN_OVER_IP              Yes  16   Unused never       -
PORT_ACTIVATION_PKG           No    0   Unused             -
SME_FOR_IPS_184_PKG           No    0   Unused             Grace 86D 5H
SAN_EXTN_OVER_IP_18_4         No    0   Unused             -
SAN_EXTN_OVER_IP_IPS2         Yes   1   Unused never       1 license(s) missing
SAN_EXTN_OVER_IP_IPS4         No    0   Unused             -
10G_PORT_ACTIVATION_PKG       No    0   Unused             -
SAN_EXTN_OVER_IP_18_4        No    0   Unused             -
STORAGE_SERVICES_ENABLER_PKG  Yes   1   Unused never       1 license(s) missing
--------------------------------------------------------------------------------
**** WARNING: License file(s) missing. ****
ips-hac1#
```

License Transfers Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.

Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Displaying License Information

Use the **show license** commands to display all license information configured on this switch (see Examples 3-2 to 3-7).

Example 3-2 Displays Information About Current License Usage

```
switch# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 No    -   Unused             Grace 79D 16H
MAINFRAME_PKG                 No    -   Unused             Grace expired
ENTERPRISE_PKG                Yes   -   Unused never       license missing
DMM_FOR_SSM_PKG               No    0   Unused             -
SAN_EXTN_OVER_IP              Yes  16   Unused never       -
PORT_ACTIVATION_PKG           No    0   Unused             -
SME_FOR_IPS_184_PKG           No    0   Unused             Grace 86D 5H
SAN_EXTN_OVER_IP_18_4         No    0   Unused             -
SAN_EXTN_OVER_IP_IPS2         Yes   1   Unused never       1 license(s) missing
SAN_EXTN_OVER_IP_IPS4         No    0   Unused             -
10G_PORT_ACTIVATION_PKG       No    0   Unused             -
SAN_EXTN_OVER_IP_18_4       No    0   Unused             -
STORAGE_SERVICES_ENABLER_PKG  Yes   1   Unused never       1 license(s) missing
--------------------------------------------------------------------------------
```

Example 3-3 Displays the List of Features in a Specified Package

```
switch# show license usage ENTERPRISE_PKG
Application
-----------
ivr
qos_manager
-----------
```

Example 3-4 Displays the Host ID for the License

```
switch# show license host-id
License hostid: FOX0646S017
```

Note Use the entire ID that appears after the colon (:) sign.

Example 3-5 Displays All Installed License Key Files and Their Contents

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
        HOSTID=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
Evaluation.lic:
SERVER this_host ANY
VENDOR cisco
```

```
INCREMENT MAINFRAME_PKG cisco 1.0 30-Dec-2003 uncounted \
        HOSTID=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Example 3-6 Displays a List of Installed License Key Files

```
switch# show license brief
Enterprise.lic
Ficon.lic
FCIP.lic
```

Example 3-7 Displays the Contents of a Specified License Key File

```
switch# show license file Permanent.lic
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
        HOSTID=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

CHAPTER 4

On-Demand Port Activation Licensing

This chapter describes how to use the on-demand port activation licensing feature on the Cisco MDS 9124 Fabric Switch, the Cisco MDS 9134 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter. This chapter contains the following sections:

About On-Demand Port Activation Licensing

As of Cisco MDS SAN-OS Release 3.1(1), you can expand your SAN connectivity as needed by enabling users to purchase and install additional port licenses. By default, all ports are eligible for license activation. On the Cisco MDS 9124 Fabric Switch, licenses are allocated sequentially. However, you can move or reassign licenses to any eligible port on the switch.

On the Cisco MDS 9134 Fabric Switch, the first 32 ports operate at 1 Gbps, 2 Gbps, or 4 Gbps. The switch has two ports that operate at 10 Gbps. Licenses are allocated sequentially.

On the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter, licenses for internal ports are allocated as the ports come up. Licenses for external ports are allocated sequentially.

Port-Naming Conventions

Table 4-1 describes the port-naming conventions for the four Cisco Fabric switches.

Table 4-1 Port-Naming Conventions for Cisco Fabric Switches

Cisco MDS 9124 Switch	Cisco MDS 9134 Switch	Cisco Fabric Switch for HP c-Class BladeSystem	Cisco Fabric Switch for IBM BladeCenter
fc1/1 through fc1/24	fc1/1 through fc1/34	Internal ports: bay1 through bay16	Internal ports: bay1 through bay14
		External ports: ext1 through ext8	External ports: ext0 and ext15 through ext19

Port Licensing

On the Cisco MDS 9124 Switch, the first eight ports are licensed by default. You are not required to perform any tasks beyond the default configuration unless you prefer to immediately activate additional ports, make ports ineligible, or move port licenses.

Figure 4-1 shows the ports that are licensed by default for the Cisco MDS 9124 Switch.

Figure 4-1 Cisco MDS 9124 Switch Default Port Licenses (fc1/1 - fc1/8)

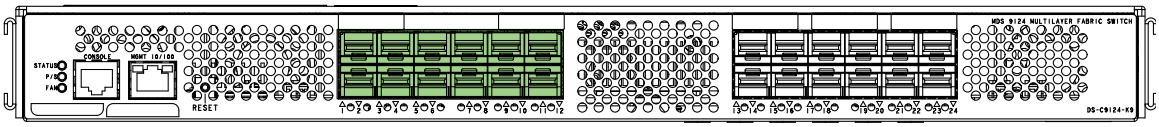

If you need additional connectivity, you can activate additional ports in 8-port increments with each on-demand port activation license, up to a total of 24 ports.

On the Cisco MDS 9134 Switch, the first 24 ports that can operate at 1 Gbps, 2 Gbps, or 4 Gbps are licensed by default. If you need additional connectivity, you can activate the remaining eight ports with one on-demand port activation license. A separate 10G license file is required to activate the remaining two 10-Gbps ports.

Figure 4-2 shows the ports that are licensed by default for the Cisco MDS 9134 Switch.

Figure 4-2 Cisco MDS 9134 Switch Default Port Licenses (fc1/1 - fc1/24)

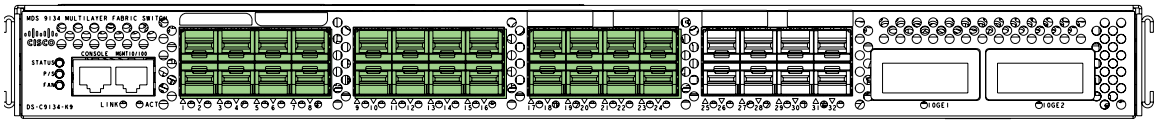

Figure 4-3 shows the external ports that are licensed by default for the Cisco Fabric Switch for HP c-Class BladeSystem.

Figure 4-3 Cisco Fabric Switch for HP c-Class BladeSystem Default Port Licenses (ext1 - ext4)

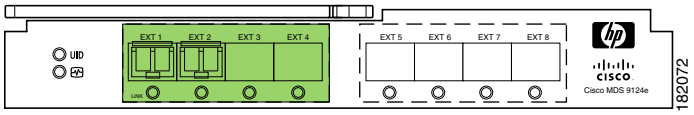

On the Cisco Fabric Switch for HP c-Class BladeSystem, any eight internal ports andthe external ports (ext1 through ext4) are licensed by default. A single on-demand port activation license is required to use the remaining eight internal and four external ports.

On the Cisco Fabric Switch for IBM BladeCenter, any seven internal ports and the external ports(ext0, ext15 and ext16) are licensed by default. A single on-demand port activation license is required to use the remaining seven internal and three external ports.

Figure 4-4 shows the external ports that are licensed by default for the Cisco Fabric Switch for IBM BladeCenter.

Figure 4-4 Cisco Fabric Switch for IBM BladeCenter Default Port Licenses (ext0, ext15 - ext16)

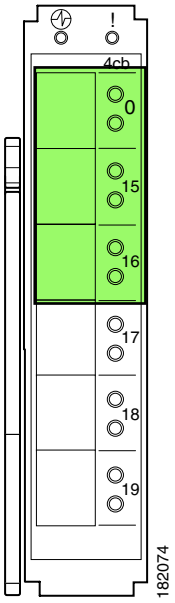

If you do not prefer to accept the default behavior and would rather assign a license to a specific port, make the port ineligible to receive a license, or move licenses among ports, refer to the "Configuring Port Activation Licenses" section on page 4-10.

Default Configuration

Example 4-1 shows the default port license configuration for the Cisco MDS 9124 Switch.

Example 4-1 Cisco MDS 9124 Switch Default Port License Configuration

```
switch# show port-license
Available port activation licenses are 0
--------------------------------------------------
  Interface    Cookie       Port Activation License
--------------------------------------------------
   fc1/1      16777216           acquired
   fc1/2      16781312           acquired
   fc1/3      16785408           acquired
   fc1/4      16789504           acquired
   fc1/5      16793600           acquired
   fc1/6      16797696           acquired
   fc1/7      16801792           acquired
   fc1/8      16805888           acquired
   fc1/9      16809984           eligible
   fc1/10     16814080           eligible
   fc1/11     16818176           eligible
   fc1/12     16822272           eligible
   fc1/13     16826368           eligible
   fc1/14     16830464           eligible
   fc1/15     16834560           eligible
   fc1/16     16838656           eligible
   fc1/17     16842752           eligible
   fc1/18     16846848           eligible
   fc1/19     16850944           eligible
   fc1/20     16855040           eligible
   fc1/21     16859136           eligible
   fc1/22     16863232           eligible
   fc1/23     16867328           eligible
   fc1/24     16871424           eligible
```

Note The cookie is used to acquire a license. Use the **show license usage PORT_ACTIVATION_PKG** command to show the cookies for acquired licenses.

```
switch# show license usage PORT_ACTIVATION_PKG
Application
-----------
16777216
16797696
16781312
16793600
16785408
16805888
16789504
16801792
-----------
```

Example 4-2 shows the default port license configuration for the Cisco MDS 9134 Switch.

Example 4-2 Cisco MDS 9134 Switch Default Port License Configuration

```
switch# show port-license
Available port activation licenses are 0
Available 10G port activation licenses are 0
```

```
-------------------------------------------------
  Interface    Cookie       Port Activation License
-------------------------------------------------
   fc1/1      16777216           acquired
   fc1/2      16781312           acquired
   fc1/3      16785408           acquired
   fc1/4      16789504           acquired
   fc1/5      16793600           acquired
   fc1/6      16797696           acquired
   fc1/7      16801792           acquired
   fc1/8      16805888           acquired
   fc1/9      16809984           acquired
   fc1/10     16814080           acquired
   fc1/11     16818176           acquired
   fc1/12     16822272           acquired
   fc1/13     16826368           acquired
   fc1/14     16830464           acquired
   fc1/15     16834560           acquired
   fc1/16     16838656           acquired
   fc1/17     16842752           acquired
   fc1/18     16846848           acquired
   fc1/19     16850944           acquired
   fc1/20     16855040           acquired
   fc1/21     16859136           acquired
   fc1/22     16863232           acquired
   fc1/23     16867328           acquired
   fc1/24     16871424           acquired
   fc1/25     16875520           eligible
   fc1/26     16879616           eligible
   fc1/27     16883712           eligible
   fc1/28     16887808           eligible
   fc1/29     16891904           eligible
   fc1/30     16896000           eligible
   fc1/31     16900096           eligible
   fc1/32     16904192           eligible
   fc1/33     16908288           eligible
   fc1/34     16912384           eligible
```

Note The cookie is used to acquire a license. Use the **show license usage PORT_ACTIVATION_PKG** command to show the cookies for acquired licenses.

```
switch# show license usage PORT_ACTIVATION_PKG
Application
-----------
16777216
16797696
16781312
16793600
16785408
16805888
16789504
16801792
16809984
16859136
16814080
16826368
16838656
16834560
16842752
16818176
16822272
```

```
16830464
16855040
16850944
16846848
16867328
16871424
16863232
-----------
```

Example 4-3 shows the default port license configuration for the Cisco Fabric Switch for HP c-Class BladeSystem.

Note The first eight internal ports that come up acquire a license.

Example 4-3 Cisco Fabric Switch for HP c-Class BladeSystem Default Port License Configuration

```
switch# show port-license
 Available ext port activation licenses are 0
 Available bay port activation licenses are 0
--------------------------------------------------
  Interface    Cookie       Port Activation License
--------------------------------------------------
   bay1       16838656           acquired
   bay2       16834560           eligible
   bay3       16818176           acquired
   bay4       16809984           eligible
   bay5       16789504           acquired
   bay6       16781312           eligible
   bay7       16805888           eligible
   bay8       16863232           acquired
   bay9       16850944           acquired
   bay10      16842752           acquired
   bay11      16822272           acquired
   bay12      16826368           eligible
   bay13      16785408           acquired
   bay14      16797696           eligible
   bay15      16801792           eligible
   bay16      16859136           eligible
   ext1       16814080           acquired
   ext2       16830464           acquired
   ext3       16846848           acquired
   ext4       16855040           acquired
   ext5       16871424           eligible
   ext6       16867328           eligible
   ext7       16793600           eligible
   ext8       16777216           eligible
```

Note The cookie is used to acquire a license. Use the **show license usage PORT_ACTIVATION_PKG** command to show the cookies for acquired licenses.

```
switch# show license usage PORT_ACTIVATION_PKG
Application
-----------
16785408
16789504
16793600
16814080
16818176
```

```
16822272
16838656
16842752
16850944
16863232
16867328
16855040
-----------
```

Example 4-4 shows the default port license configuration for the Cisco Fabric Switch for IBM BladeCenter.

Note The first seven internal ports that come up acquire a license.

Example 4-4 Cisco Fabric Switch for IBM BladeCenter Default Port License Configuration

```
switch# show port-license
Available ext port activation licenses are 0
Available bay port activation licenses are 0
--------------------------------------------------
  Interface    Cookie       Port Activation License
--------------------------------------------------
   bay1       16850944           eligible
   bay2       16838656           eligible
   bay3       16842752           acquired
   bay4       16834560           eligible
   bay5       16822272           acquired
   bay6       16818176           eligible
   bay7       16826368           acquired
   bay8       16809984           eligible
   bay9       16797696           acquired
   bay10      16781312           eligible
   bay11      16785408           acquired
   bay12      16789504           eligible
   bay13      16801792           acquired
   bay14      16805888           acquired
   ext0       16846848           acquired
   ext15      16855040           acquired
   ext16      16830464           acquired
   ext17      16814080           eligible
   ext18      16793600           eligible
   ext19      16777216           eligible
```

Note The cookie is used to acquire a license. Use the **show license usage PORT_ACTIVATION_PKG** command to show the cookies for acquired licenses.

```
switch# show license usage PORT_ACTIVATION_PKG
Application
-----------
16830464
16826368
16818176
16822272
16834560
16838656
16846848
16850944
16855040
```

```
16842752
-----------
```

License Status Definitions

Table 4-2 defines the port activation license status terms.

Table 4-2 Port Activation License Status Definitions

Port Activation License Status	Definition
acquired	The port is licensed and active.
eligible	The port is eligible to receive a license but does not yet have one. See Chapter 3, "Obtaining and Installing Licenses," for information about how to obtain and install the PORT_ACTIVATION_PKG and license key file.
ineligible	The port is not allowed to receive a license.

By default, when you install additional port license activation packages, the activation status of ports changes from "eligible" to "acquired." If you prefer to accept the default behavior, no further action is required.

Note You can uninstall licenses for ports not in use; however, you cannot uninstall default licenses.

Table 4-3 describes the port license assignments for the Cisco MDS 9124 Switch.

Table 4-3 Default Port License Assignments for Cisco MDS 9124 Switch

License Package (PORT_ACTIVATION_PKG)	Assigned to Ports on the Cisco MDS 9124 Switch
Default	1-8
First PORT_ACTIVATION_PKG	9-16
Second PORT_ACTIVATION_PKG	17-24

You can use the **show license usage** command to view any licenses assigned to a switch. If a license is in use, the status displayed is "In use." If a license is installed but no ports have acquired a license, then the status displayed is "Unused."

The default license package for the Cisco MDS 9124 Switch is as follows:

```
switch# show license usage

Feature                      Ins       Lic   Status     Expiry Date  Comments
                                       Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 Yes       -    Unused        never       -
ENTERPRISE_PKG                Yes       -    In use        never       -
PORT_ACTIVATION_PKG           No        8    In use        never       -
```

```
10G_PORT_ACTIVATION_PKG       No       0    Unused                    -
--------------------------------------------------------------------------------
```

Note The PORT_ACTIVATION_PKG does not appear as installed if you have only the default license installed.

Table 4-4 describes the port license assignments for the Cisco MDS 9134 Switch.

Table 4-4 Default Port License Assignments for Cisco MDS 9134 Switch

License Package (PORT_ACTIVATION_PKG)	Assigned to Ports on the Cisco MDS 9134 Switch
Default	1-24
PORT_ACTIVATION_PKG	25-32
10G_PORT_ACTIVATION_PKG	33-34

You can use the **show license usage** command to view any licenses assigned to a switch. If a license is in use, the status displayed is "In use." If a license is installed but no ports have acquired a license, then the status displayed is "Unused."

The default license package for the Cisco MDS 9134 Switch is as follows:

```
switch# show license usage

Feature                      Ins      Lic   Status     Expiry Date  Comments
                                      Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 Yes      -    Unused        never       -
ENTERPRISE_PKG                Yes      -    In use        never       -
PORT_ACTIVATION_PKG           No       24   In use        never       -
10G_PORT_ACTIVATION_PKG       yes      2    Unused        never       -
--------------------------------------------------------------------------------
```

Note The PORT_ACTIVATION_PKG does not appear as installed if you have only the default license installed.

Table 4-5 describes the port license assignments for the Cisco Fabric Switch for HP c-Class BladeSystem.

Table 4-5 Default Port License Assignments for Cisco Fabric Switch for HP c-Class BladeSystem

License Package (PORT_ACTIVATION_PKG)	Assigned to Ports on the Cisco Fabric Switch for HP c-Class BladeSystem
Default	Any eight internal ports and the four external ports ext1 through ext4.
PORT_ACTIVATION_PKG	A single license required for the remaining eight internal and four external ports.

You can use the **show license usage** command to view any licenses assigned to a switch. The default license package for the Cisco Fabric Switch for HP c-Class BladeSystem is as follows:

```
switch# show license usage

Feature                      Ins      Lic    Status     Expiry Date  Comments
                                      Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 No       -     Unused                   -
ENTERPRISE_PKG                No       -     Unused                   -
PORT_ACTIVATION_PKG           No      12     In use          never    -
10G_PORT_ACTIVATION_PKG       No       0     Unused                   -
--------------------------------------------------------------------------------
```

Note The PORT_ACTIVATION_PKG does not appear as installed if you have only the default license installed.

Table 4-6 describes the port license assignments for the Cisco Fabric Switch for IBM BladeCenter.

Table 4-6 Default Port License Assignments for Cisco Fabric Switch for IBM BladeCenter

License Package (PORT_ACTIVATION_PKG)	Assigned to Ports on the Cisco Fabric Switch for IBM BladeCenter
Default	Any seven internal ports and the three external ports ext0, ext15 and ext16.
PORT_ACTIVATION_PKG	A single license required for the remaining seven internal and three external ports.

You can use the **show license usage** command to view any licenses assigned to a switch. The default license package for the Cisco Fabric Switch for IBM BladeCenter is as follows:

```
switch# show license usage

Feature                      Ins      Lic    Status     Expiry Date  Comments
                                      Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                 No       -     Unused                   -
ENTERPRISE_PKG                No       -     Unused                   -
PORT_ACTIVATION_PKG           No      10     In use          never    -
10G_PORT_ACTIVATION_PKG       No       0     Unused                   -
--------------------------------------------------------------------------------
```

Note The PORT_ACTIVATION_PKG does not appear as installed if you have only the default license installed.

Configuring Port Activation Licenses

This section contains the following topics:

Making a Port Eligible for a License

By default, all ports are eligible to receive a license. However, if a port has already been made ineligible and you prefer to activate it, then you must make that port eligible by using the **port-license** command.

To make a port eligible to acquire a license, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Specifies the port interface that you want to make eligible for a license. **Note** The name of the port depends on the switch you are using. See "Port-Naming Conventions" section on page 4-2 for information on port names.
Step 3	`switch(config-if)# port-license`	Makes the port eligible to acquire a license.
	`switch(config-if)# no port-license`	Removes a license from a port if it already has been assigned, and also makes the port ineligible to acquire a license. **Note** You can remove licenses only from ports that are not in an administrative shutdown state.

Acquiring a License for a Port

If you do not prefer to accept the default on-demand port license assignments, you will need to first acquire licenses for ports to which you want to move the license.

To acquire a license for a port, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Specifies the port interface for which you want to acquire a license. **Note** The name of the port depends on the switch you are using. See "Port-Naming Conventions" section on page 4-2 for information on port names.
Step 3	`switch(config-if)# port-license acquire`	Grants a license to a port or range of ports.
	`switch(config-if)# no port-license`	Removes a license from a port or range of ports.

Moving Licenses Among Ports

Note On the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter, you can only move the license for internal ports among internal ports. Licenses from an internal port cannot be moved to an external port and vice versa. Licenses for external ports can only be moved among external ports.

You can move a license from a port (or range of ports) at any time. If you attempt to move a license to a port and no license is available, then the switch returns the message "port activation license not available."

Note Once internal ports are licensed on the Cisco Fabric Switch for HP c-Class BladeSystem or the Cisco Fabric Switch for IBM BladeCenter, if the user enters the **copy running-config startup-config** command, then on the next reload, these ports will retain the licenses.

To move a license from one port to another (in this example, from fc1/1 to fc1/24), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Specifies the port interface from which you want to move a license. **Note** The name of the port depends on the switch you are using. See the "Port-Naming Conventions" section on page 4-2 for information on port names.
Step 3	`switch(config-if)# no port-license`	Removes the license from port fc1/1 and makes the port ineligible to acquire a license. **Note** The port needs to be shut down for this command to take effect.
Step 4	`switch(config-if)# exit` `switch(config)#`	Exits interface mode for fc1/1.
Step 5	`switch(config)# interface fc1/24` `switch(config-if)#`	Specifies the port interface to which you want to move the license. **Note** The name of the port depends on the switch you are using. See "Port-Naming Conventions" section on page 4-2 for information on port names.
Step 6	`switch(config-if)# port-license acquire`	Grants a license to port fc1/24. **Note** The port needs to be shut down for this command to take effect.
Step 7	`switch(config-if)# end`	Returns to EXEC mode.

On-Demand Port Activation License Example

The following example shows how to do the following tasks:

- Make a port ineligible
- Install port activation licenses
- Move licenses from one port to another

If you do not want to accept the default behavior, or you need flexibility in terms of which ports acquire a license, you may want to make a port ineligible. For example, if the first eight ports have a license, but you want to move a license from port 7 to port 9, then you would need to make a port ineligible. Or, if you have a port that should never acquire a license, you can make it ineligible and it will not be a candidate for a license when additional licenses are installed.

This example is based on the default configuration for the Cisco MDS 9124 Switch.

Step 1 Display the default port license configuration.

```
switch# show port-license
Available port activation licenses are 0
-----------------------------------------------------
  Interface    Cookie       Port Activation License
-----------------------------------------------------
   fc1/1      16777216           acquired
   fc1/2      16781312           acquired
   fc1/3      16785408           acquired
   fc1/4      16789504           acquired
   fc1/5      16793600           acquired
   fc1/6      16797696           acquired
   fc1/7      16801792           acquired
   fc1/8      16805888           acquired
   fc1/9      16809984           eligible
   fc1/10     16814080           eligible
   fc1/11     16818176           eligible
   fc1/12     16822272           eligible
   fc1/13     16826368           eligible
   fc1/14     16830464           eligible
   fc1/15     16834560           eligible
   fc1/16     16838656           eligible
   fc1/17     16842752           eligible
   fc1/18     16846848           eligible
   fc1/19     16850944           eligible
   fc1/20     16855040           eligible
   fc1/21     16859136           eligible
   fc1/22     16863232           eligible
   fc1/23     16867328           eligible
   fc1/24     16871424           eligible
```

Step 2 Install an additional license package. See Chapter 3, “Obtaining and Installing Licenses,” for information about how to obtain and install the PORT_ACTIVATION_PKG and license key file.

```
switch# install license bootflash:license_file.lic
Installing license ..done
```

Note If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the filename specified in the license key file is used to install the license.

Step 3 Make port fc1/8 ineligible to receive a license.

Note When you make a port ineligible, the license does not automatically transfer to another port.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc1/8
switch(config-if)# no port-license
switch(config-if)# end
switch# show port-license
Available port activation licenses are 1
-----------------------------------------------------
  Interface    Cookie       Port Activation License
-----------------------------------------------------
   fc1/1      16777216           acquired
   fc1/2      16781312           acquired
   fc1/3      16785408           acquired
   fc1/4      16789504           acquired
   fc1/5      16793600           acquired
   fc1/6      16797696           acquired
   fc1/7      16801792           acquired
   fc1/8      16805888           ineligible
   fc1/9      16809984           eligible
   fc1/10     16814080           eligible
```

Step 4 Display the licensed features to confirm that you have successfully installed PORT_ACTIVATION_PKG.

```
switch# show license default
Feature                                 Default License Count
--------------------------------------------------------------------------------
FM_SERVER_PKG                                        -
ENTERPRISE_PKG                                       -
PORT_ACTIVATION_PKG                                  8
10G_PORT_ACTIVATION_PKG                              0
switch# ---------------------------------------------------------------------
```

Step 5 Display the port license configuration to confirm that additional ports have acquired a license.

Note Port fc1/8 remains ineligible and one license remains available. Ports fc1/9 through fc1//16 have acquired an additional license.

```
switch# show port-license
Available port activation licenses are 1
-----------------------------------------------------
  Interface    Cookie       Port Activation License
-----------------------------------------------------
   fc1/1      16777216           acquired
   fc1/2      16781312           acquired
   fc1/3      16785408           acquired
   fc1/4      16789504           acquired
   fc1/5      16793600           acquired
   fc1/6      16797696           acquired
   fc1/7      16801792           acquired
   fc1/8      16805888           ineligible
   fc1/9      16809984           acquired
   fc1/10     16814080           acquired
   fc1/11     16818176           acquired
   fc1/12     16822272           acquired
   fc1/13     16826368           acquired
   fc1/14     16830464           acquired
```

```
   fc1/15     16834560          acquired
   fc1/16     16838656          acquired
   fc1/17     16842752          eligible
   fc1/18     16846848          eligible
   fc1/19     16850944          eligible
   fc1/20     16855040          eligible
   fc1/21     16859136          eligible
   fc1/22     16863232          eligible
   fc1/23     16867328          eligible
   fc1/24     16871424          eligible
```

Step 6 Move the remaining license to port fc1/17.

```
switch# config t
switch(config)# interface fc1/17
switch(config-int)# port-license acquire
```

Step 7 Display the port license configuration to confirm that port fc1/17 has acquired a license.

```
switch# show port-license
Available port activation licenses are 0
--------------------------------------------------
  Interface    Cookie       Port Activation License
--------------------------------------------------
   fc1/1      16777216          acquired
   fc1/2      16781312          acquired
   fc1/3      16785408          acquired
   fc1/4      16789504          acquired
   fc1/5      16793600          acquired
   fc1/6      16797696          acquired
   fc1/7      16801792          acquired
   fc1/8      16805888          ineligible
   fc1/9      16809984          acquired
   fc1/10     16814080          acquired
   fc1/11     16818176          acquired
   fc1/12     16822272          acquired
   fc1/13     16826368          acquired
   fc1/14     16830464          acquired
   fc1/15     16834560          acquired
   fc1/16     16838656          acquired
   fc1/17     16842752          acquired
   fc1/18     16846848          eligible
   fc1/19     16850944          eligible
   fc1/20     16855040          eligible
   fc1/21     16859136          eligible
   fc1/22     16863232          eligible
   fc1/23     16867328          eligible
   fc1/24     16871424          eligible
```

Step 8 Make this configuration your startup configuration by saving the new port license configuration into nonvolatile storage. Once you complete this step, the running and the startup copies of the configuration are identical.

```
switch# copy running-config startup-config
```

Step 9 Display and/or confirm the licenses in the running configuration by entering the **show running config** command.

```
switch# show running config
...
interface fc1/1
  switchport trunk mode auto
  port-license acquire
  channel-group 122 force
```

```
  no shutdown

interface fc1/2
  switchport trunk mode auto
  port-license acquire
  channel-group 122 force
  no shutdown

interface fc1/3
  switchport trunk mode auto
  port-license acquire
  no shutdown

interface fc1/4
  port-license acquire
  no shutdown

interface fc1/5
  switchport trunk mode auto
  port-license acquire
  port-track interface fc1/13
  port-track interface fc1/21
  port-track interface fc1/24
  port-track interface port-channel 122
  no shutdown

interface fc1/6
  switchport trunk mode off
  port-license acquire
  fcsp auto-active
  no shutdown
```

CHAPTER 5

Initial Configuration

This chapter includes the following sections:

Starting a Switch in the Cisco MDS 9000 Family

The following procedure is a review of the tasks you should have completed during hardware installation, including starting up the switch. These tasks must be completed before you can configure the switch.

Before you can configure a switch, follow these steps:

Step 1 Verify the following physical connections for the new Cisco MDS 9000 Family switch:

- The console port is physically connected to a computer terminal (or terminal server).
- The management 10/100/1000 Ethernet port (mgmt0) is connected to an external hub, switch, or router.

Refer to the *Cisco MDS 9000 Family Hardware Installation Guide* (for the required product) for more information.

Tip Save the host ID information for future use (for example, to enable licensed features). The host ID information is provided in the Proof of Purchase document that accompanies the switch.

Step 2 Verify that the default console port parameters are identical to those of the computer terminal (or terminal server) attached to the switch console port:

- 9600 baud
- 8 data bits
- 1 stop bit
- No parity

Note On Cisco terminal servers, issue the following commands starting in EXEC mode:

```
switch# config t
switch(config)# line 1
switch(config)# no flush-at-activation
switch(config)# line 1
switch(config)# exit
switch# copy running-config startup-config
```

This configuration ensures that the MDS switch does not receive random characters that might cause it to hang.

Step 3 Power on the switch. The switch boots automatically and the `switch#` prompt appears in your terminal window.

Initial Setup Routine

The first time that you access a switch in the Cisco MDS 9000 Family, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the supervisor module Ethernet interface. This information is required to configure and manage the switch.

Note The IP address can only be configured from the CLI. When you power up the switch for the first time assign the IP address. After you perform this step, the Cisco MDS 9000 Family Fabric Manager can reach the switch through the console port.

Preparing to Configure the Switch

Before you configure a switch in the Cisco MDS 9000 Family for the first time, you need the following information:

- Administrator password, including:
  - Creating a password for the administrator (required).
  - Creating an additional login account and password (optional).

Note If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password. You must configure a password that meets the requirements listed in the "Characteristics of Strong Passwords" section on page 31-12.

- IPv4 address or IPv6 address for the switch management interface—The management interface can be an out-of-band Ethernet interface or an in-band Fibre Channel interface (recommended).
- If you are using an IPv4 address for the management interface, you need the following information:
  - IPv4 subnet mask for the switch's management interface (optional).
  - Destination IPv4 prefix, destination IPv4 prefix subnet mask, and next hop IPv4 address, if you want to enable IP routing.
  - IPv4 address of the default gateway (optional).
- SSH service on the switch—To enable this optional service, select the type of SSH key (dsa/rsa/rsa1) and number of key bits (768 to 2048).
- DNS IPv4 address or IPv6 address (optional).
- Default domain name (optional).
- NTP server IPv4 address or IPv6 address (optional).
- SNMP community string (optional).
- Switch name—This is your switch prompt (optional).

Note If you are using IPv4, be sure to configure the IPv4 route, the IPv4 default network address, and the IPv4 default gateway address to enable SNMP access. If IP routing is enabled, the switch uses the IPv4 route and the default network IPv4 address. If IP routing is disabled, the switch uses the default gateway IPv4 address.

Default Login

All Cisco MDS 9000 Family switches have the network administrator as a default user (admin). You cannot change the default user at any time (see the "Role-Based Authorization" section on page 37-1).

There is no default password so you must explicitly configure a strong password. If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password (see the). If you configure and subsequently forget this new password, you have the option to recover this password (see the "Recovering the Administrator Password" section on page 37-19).

Note The Cisco Fabric Switch for IBM BladeCenter does not use admin as the default user. Rather, the default user is USERID because there is no console access to the switch. You cannot delete the user USERID on this switch. The password for this default user is PASSW0RD, where the "0" is a zero. You can change this password; however, a write erase operation restores the default password. There is no initial setup menu.

Also note that you should not bring up the loader> prompt; the only way to fix this condition is to RMA the switch.

The following commands are not allowed on the Cisco Fabric Switch for IBM BladeCenter: **write erase boot** and **init system**; nor can you boot variables manually.

Note If you issue a **write erase** command and reload the switch, you must reconfigure the default user (admin) password using the setup procedure.

Setup Options

The setup scenario differs based on the subnet to which you are adding the new switch. You must configure a Cisco MDS 9000 Family switch with an IP address to enable management connections from outside of the switch.

Note Some concepts such as out-of-band management and in-band management are briefly explained here. These concepts are explained in more detail in subsequent chapters.

- Out-of-band management—This feature provides a connection to the network through a supervisor module front panel Ethernet port (see Figure 5-1).
- In-band management—This feature provides IP over Fibre Channel (IPFC) to manage the switches. The in-band management feature is transparent to the network management system (NMS). Instead of conventional Ethernet physical media, switches in the Cisco MDS 9000 Family use IPFC as the transport mechanism (see Figure 5-1 and Chapter 43, "Configuring IP Services").

Figure 5-1 Management Access to Switches

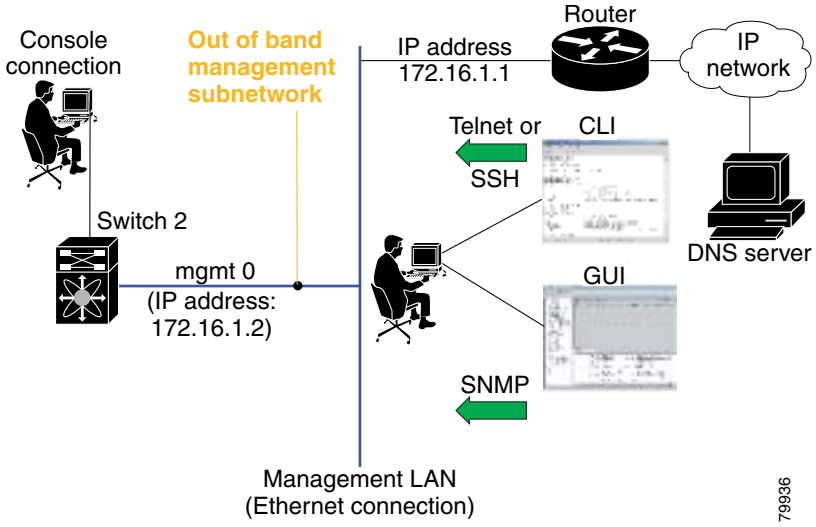

Assigning Setup Information

This section describes how to initially configure the switch for both out-of-band and in-band management.

Note Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point. Entering the new password for the administrator is a requirement and cannot be skipped. See the "Characteristics of Strong Passwords" section on page 31-12.

Tip If you do not wish to answer a previously configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

Note The setup script only supports IPv4 for the management interface. For information on configuring IPv6 on the management interface, see the Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces."

Configuring Out-of-Band Management

Note You can configure both in-band and out-of-band configuration together by entering **Yes** in both Step 11c and Step 11d in the following procedure.

To configure the switch for first time out-of-band access, follow these steps:

Step 1 Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

```
Enter the password for admin: 2004AsdfLkjh18
```

Tip If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. You must explicitly configure a password that meets the requirements listed in the "Characteristics of Strong Passwords" section on page 31-12.

Step 3 Enter **yes** to enter the setup mode.

```
This setup utility will guide you through the basic configuration of the system. Setup
configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. MDS devices must be
registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away
remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter the new password for the administrator (admin is the default).

```
Enter the password for admin: admin
```

Step 5 Enter **yes** (no is the default) to create additional accounts.

```
Create another login account (yes/no) [n]: yes
```

While configuring your initial setup, you can create an additional user account (in the network-admin role) besides the administrator's account. See the "Role-Based Authorization" section on page 37-1 for information on default roles and permissions.

Note User login IDs must contain non-numeric characters.

a. Enter the user login ID.

```
Enter the user login ID: user_name
```

b. Enter the user password.

```
Enter the password for user_name: user-password
```

Step 6 Enter **yes** (yes is the default) to create an SNMPv3 account.

```
Configure SNMPv3 Management parameters (yes/no) [y]: yes
```

a. Enter the user name (admin is the default).

```
SNMPv3 user name [admin]: admin
```

b. Enter the SNMPv3 password (minimum of eight characters). The default is **admin123**.

```
SNMPv3 user authentication password: admin_pass
```

Note By default, if the admin password is at least eight characters, then the SNMP authentication password is the same as the admin password (at least eight characters). If the admin password is less than eight characters, then you need to provide a new password for SNMP. The admin password can have a minimum of one character, but the SNMP authentication password must have a minimum of eight characters.

Step 7 Enter **yes** (no is the default) to configure the read-only or read-write SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: yes
```

a. Enter the SNMP community string.

```
SNMP community string: snmp_community
```

Step 8 Enter a name for the switch.

Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

```
Enter the switch name: switch_name
```

Step 9 Enter **yes** (yes is the default) to configure out-of-band management.

```
Continue with Out-of-band (mgmt0) management configuration? [yes/no]: yes
```

a. Enter the mgmt0 IPv4 address.

```
Mgmt0 IPv4 address: ip_address
```

b. Enter the mgmt0 IPv4 subnet mask.

```
Mgmt0 IPv4 netmask: subnet_mask
```

Step 10 Enter **yes** (yes is the default) to configure the IPv4 default gateway (recommended).

```
Configure the default-gateway: (yes/no) [y]: yes
```

a. Enter the default gateway IPv4 address.

```
IPv4 address of the default-gateway: default_gateway
```

Step 11 Enter **yes** (**no** is the default) to configure advanced IP options such as in-band management, static routes, default network, DNS, and domain name.

```
Configure Advanced IP options (yes/no)? [n]: yes
```

a. Enter **no** (no is the default) at the in-band management configuration prompt.

```
Continue with in-band (VSAN1) management configuration? (yes/no) [no]: no
```

b. Enter **yes** (yes is the default) to enable IPv4 routing capabilities.

```
Enable the ip routing? (yes/no) [y]: yes
```

c. Enter **yes** (yes is the default) to configure a static route (recommended).

```
Configure static route: (yes/no) [y]: yes
```

Enter the destination prefix.

```
Destination prefix: dest_prefix
```

Type the destination prefix mask.

```
Destination prefix mask: dest_mask
```

Type the next hop IP address.

```
Next hop ip address: next_hop_address
```

Note Be sure to configure the IP route, the default network IP address, and the default gateway IP address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

d. Enter **yes** (yes is the default) to configure the default network (recommended).

```
Configure the default network: (yes/no) [y]: yes
```

Enter the default network IPv4 address.

Note The default network IPv4 address is the destination prefix provided in Step 11c .

```
Default network IP address [dest_prefix]: dest_prefix
```

e. Enter **yes** (yes is the default) to configure the DNS IPv4 address.

```
Configure the DNS IP address? (yes/no) [y]: yes
```

Enter the DNS IP address.

```
DNS IP address: name_server
```

f. Enter **yes** (default is no) to configure the default domain name.

```
Configure the default domain name? (yes/no) [n]: yes
```

Enter the default domain name.

```
Default domain name: domain_name
```

Step 12 Enter **yes** (yes is the default) to enable the Telnet service.

```
Enable the telnet service? (yes/no) [y]: yes
```

Step 13 Enter **yes** (no is the default) to enable the SSH service.

```
Enabled SSH service? (yes/no) [n]: yes
```

Step 14 Enter the SSH key type (see the "Overwriting a Generated Key-Pair" section on page 31-17) that you would like to generate.

```
Type the SSH key you would like to generate (dsa/rsa/rsa1)? dsa
```

Step 15 Enter the number of key bits within the specified range.

```
Enter the number of key bits? (768 to 2048): 768
```

Step 16 Enter **yes** (no is the default) to configure the NTP server.

```
Configure NTP server? (yes/no) [n]: yes
```

a. Enter the NTP server IPv4 address.

```
NTP server IP address: ntp_server_IP_address
```

Step 17 Enter **shut** (shut is the default) to configure the default switch port interface to the shut (disabled) state.

```
Configure default switchport interface state (shut/noshut) [shut]: shut
```

Note The management ethernet interface is not shut down at this point—only the Fibre Channel, iSCSI, FCIP, and Gigabit Ethernet interfaces are shut down.

Step 18 Enter **on** (on is the default) to configure the switch port trunk mode.

```
Configure default switchport trunk mode (on/off/auto) [on]: on
```

Step 19 Enter **yes** (yes is the default) to configure the switchport mode F.

```
Configure default switchport mode F (yes/no) [n]: y
```

Step 20 Enter **on** (off is the default) to configure the PortChannel auto-create state.

```
Configure default port-channel auto-create state (on/off) [off]: on
```

Step 21 Enter **permit** (deny is the default) to deny a default zone policy configuration.

```
Configure default zone policy (permit/deny) [deny]: permit
```

Permits traffic flow to all members of the default zone.

Note If you are executing the setup script after issuing a **write erase** command, you explicitly must change the default zone policy to permit for VSAN 1 after finishing the script using the following commands:

```
switch# config t
switch(config)# zone default-zone permit vsan 1
```

Step 22 Enter **yes** (no is the default) to enable a full zone set distribution (see the "Zone Set Distribution" section on page 23-13).

```
Enable full zoneset distribution (yes/no) [n]: yes
```

Overrides the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Step 23 Enter **no** (no is the default) if you are satisfied with the configuration.

```
The following configuration will be applied:
  username admin password admin_pass role network-admin
  username user_name password user_pass role network-admin
  snmp-server community snmp_community ro
```

```
  switchname switch
  interface mgmt0
    ip address ip_address subnet_mask
    no shutdown
  ip routing
  ip route dest_prefix dest_mask dest_address
  ip default-network dest_prefix
  ip default-gateway default_gateway
  ip name-server name_server
  ip domain-name domain_name
  telnet server enable
  ssh key dsa 768 force
  ssh server enable
  ntp server ipaddr ntp_server
  system default switchport shutdown
  system default switchport trunk mode on
  system default switchport mode F
  system default port-channel auto-create
  zone default-zone permit vsan 1-4093
  zoneset distribute full vsan 1-4093
Would you like to edit the configuration? (yes/no) [n]: no
```

Step 24 Enter **yes** (yes is default) to use and save this configuration:

```
Use this configuration and save it? (yes/no) [y]: yes
```

Caution If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** to save the new configuration. This ensures that the kickstart and system images are also automatically configured (see Chapter 7, "Software Images").

Configuring In-Band Management

The in-band management logical interface is VSAN 1. This management interface uses the Fibre Channel infrastructure to transport IP traffic. An interface for VSAN 1 is created on every switch in the fabric. Each switch should have its VSAN 1 interface configured with either an IPv4 address or an IPv6 address in the same subnetwork. A default route that points to the switch providing access to the IP network should be configured on every switch in the Fibre Channel fabric (see Chapter 19, "Configuring and Managing VSANs").

Note You can configure both in-band and out-of-band configuration together by entering **Yes** in both Step 9c and Step 9d in the following procedure.

To configure a switch for first time in-band access, follow these steps:

Step 1 Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

```
Enter the password for admin: 2004asdf*lkjh18
```

Tip If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. You must explicitly configure a password that meets the requirements listed in the "Characteristics of Strong Passwords" section on page 31-12.

Step 3 Enter **yes** to enter the setup mode.

```
This setup utility will guide you through the basic configuration of the system. Setup
configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. MDS devices must be
registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away
remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter **no** (no is the default) if you do not wish to create additional accounts.

```
Create another login account (yes/no) [no]: no
```

Step 5 Configure the read-only or read-write SNMP community string.

a. Enter **no** (no is the default) to avoid configuring the read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: no
```

b. Enter **no** (no is the default) to configure the read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: yes
```

c. Enter the SNMP community string.

```
SNMP community string: snmp_community
```

Step 6 Enter a name for the switch.

Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

```
Enter the switch name: switch_name
```

Step 7 Enter **no** (yes is the default) at the configuration prompt to configure out-of-band management.

```
Continue with Out-of-band (mgmt0) management configuration? [yes/no]: no
```

Step 8 Enter **yes** (yes is the default) to configure the default gateway.

```
Configure the default-gateway: (yes/no) [y]: yes
```

a. Enter the default gateway IP address.

```
IP address of the default gateway: default_gateway
```

Step 9 Enter **yes** (**no** is the default) to configure advanced IP options such as in-band management, static routes, default network, DNS, and domain name.

```
Configure Advanced IP options (yes/no)? [n]: yes
```

a. Enter **yes** (no is the default) at the in-band management configuration prompt.

```
Continue with in-band (VSAN1) management configuration? (yes/no) [no]: yes
```

Enter the VSAN 1 IPv4 address.

```
VSAN1 IPv4 address: ip_address
```

Enter the IPv4 subnet mask.

```
VSAN1 IPv4 net mask: subnet_mask
```

b. Enter **no** (yes is the default) to enable IPv4 routing capabilities.

```
Enable ip routing capabilities? (yes/no) [y]: no
```

c. Enter **no** (yes is the default) to configure a static route.

```
Configure static route: (yes/no) [y]: no
```

d. Enter **no** (yes is the default) to configure the default network.

```
Configure the default-network: (yes/no) [y]: no
```

e. Enter **no** (yes is the default) to configure the DNS IPv4 address.

```
Configure the DNS IP address? (yes/no) [y]: no
```

f. Enter **no** (no is the default) to skip the default domain name configuration.

```
Configure the default domain name? (yes/no) [n]: no
```

Step 10 Enter **no** (yes is the default) to disable the Telnet service.

```
Enable the telnet service? (yes/no) [y]: no
```

Step 11 Enter **yes** (no is the default) to enable the SSH service.

```
Enabled SSH service? (yes/no) [n]: yes
```

Step 12 Enter the SSH key type that you would like to generate.

```
Type the SSH key you would like to generate (dsa/rsa/rsa1)? rsa
```

Step 13 Enter the number of key bits within the specified range.

```
Enter the number of key bits? (768 to 1024): 1024
```

Step 14 Enter **no** (no is the default) to configure the NTP server.

```
Configure NTP server? (yes/no) [n]: no
```

Step 15 Enter **shut** (shut is the default) to configure the default switch port interface to the shut (disabled) state.

```
Configure default switchport interface state (shut/noshut) [shut]: shut
```

Note The management Ethernet interface is not shut down at this point—only the Fibre Channel, iSCSI, FCIP, and Gigabit Ethernet interfaces are shut down.

Step 16 Enter **auto** (off is the default) to configure the switch port trunk mode.

```
Configure default switchport trunk mode (on/off/auto) [off]: auto
```

Step 17 Enter **yes** (yes is the default) to configure the switchport mode F.

```
Configure default switchport mode F (yes/no) [n]: y
```

Step 18 Enter **off** (off is the default) to configure the PortChannel auto-create state.

```
Configure default port-channel auto-create state (on/off) [off]: off
```

Step 19 Enter **deny** (deny is the default) to deny a default zone policy configuration.

```
Configure default zone policy (permit/deny) [deny]: deny
```

Denies traffic flow to all members of the default zone.

Step 20 Enter **no** (no is the default) to disable a full zone set distribution (see the "Zone Set Distribution" section on page 23-13).

```
Enable full zoneset distribution (yes/no) [n]: no
```

Disables the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Step 21 Enter **no** (no is the default) if you are satisfied with the configuration.

```
The following configuration will be applied:
  username admin password admin_pass role network-admin
  snmp-server community snmp_community rw
  switchname switch
  interface vsan1
    ip address ip_address subnet_mask
    no shutdown
  ip default-gateway default_gateway
  no telnet server enable
  ssh key rsa 1024 force
  ssh server enable
  system default switchport shutdown
  system default switchport trunk mode auto
  system default switchport mode F
  no zone default-zone permit vsan 1-4093
  no zoneset distribute full vsan 1-4093
Would you like to edit the configuration? (yes/no) [n]: no
```

Step 22 Enter **yes** (yes is default) to use and save this configuration.

```
Use this configuration and save it? (yes/no) [y]: yes
```

Caution If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** to save the new configuration. This ensures that the kickstart and system images are also automatically configured (see Chapter 7, "Software Images").

Using the setup Command

To make changes to the initial configuration at a later time, you can issue the **setup** command in EXEC mode.

```
switch# setup
---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
*Note: setup always assumes a predefined defaults irrespective
of the current system configuration when invoked from CLI.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime
to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

The setup utility guides you through the basic configuration process.

Accessing the Switch

After initial configuration, you can access the switch in one of three ways (see Figure 5-2):

- Serial console access—You can use a serial port connection to access the CLI.
- In-band IP (IPFC) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager application.
- Out-of-band (10/100/1000 BASE-T Ethernet) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager application. Supervisor-1 modules support 10/100 BASE-T Ethernet and Supervisor-2 modules support 10/100/1000 BASE-T Ethernet.

Note To use the Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Figure 5-2 Switch Access Options

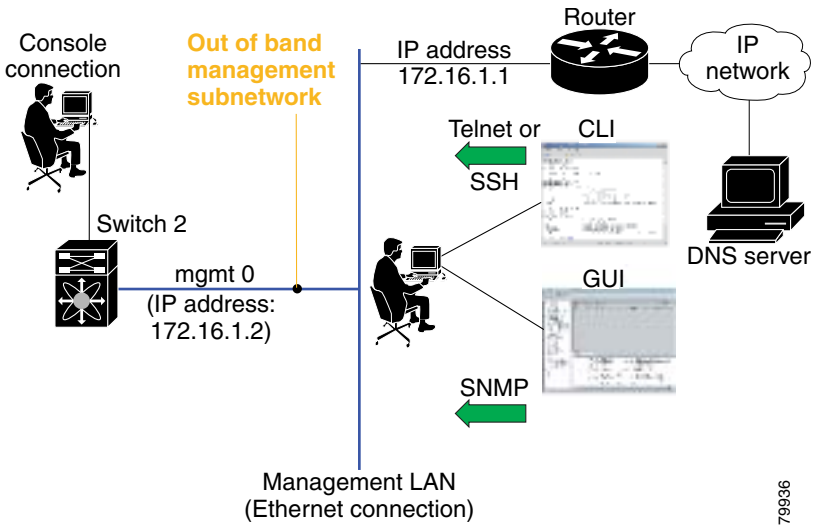

Assigning a Switch Name

Each switch in the fabric requires a unique name. You can assign names to easily identify the switch by its physical location, its SAN association, or the organization to which it is deployed. The assigned name is displayed in the command-line prompt. The switch name is limited to 20 alphanumeric characters.

Note This guide refers to a switch in the Cisco MDS 9000 Family as *switch*, and it uses the `switch#` prompt.

To change the name of the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# switchname myswitch1` `myswitch1(config)#`	Changes the switch name prompt as specified (myswitch1).
Step 3	`myswitch1(config)# no switchname` `switch(config)#`	Reverts the switch name prompt to its default (switch#).

Where Do You Go Next?

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the CLI. However, you can continue to configure other software features, or access the switch after initial configuration by using either the CLI or the Device Manager and Fabric Manager applications.

To use the Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Verifying the Module Status

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed.

To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
2    8      IP Storage Services Module       DS-X9308-SMIP      ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
8    0      Caching Services Module          DS-X9560-SMAP      ok
9    32     1/2 Gbps FC Module               DS-X9032           ok

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
2    1.3(0.106a)  0.206   20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a)  0.602   --
6    1.3(0.106a)) 0.602   --
8    1.3(0.106a)  0.702   --
9    1.3(0.106a)  0.3     22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de  JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2  JAB06350B1R
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd  JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe  JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6  JAB06280ae9

* this terminal session
```

If the status is OK or active, you can continue with your configuration (see Chapter 11, “Managing Modules”).

Configuring Date, Time, and Time Zone

Switches in the Cisco MDS 9000 Family use Universal Coordinated Time (UTC), which is the same as Greenwich Mean Time (GMT).

To change the default time on the switch, issue the **clock** command from EXEC mode.

```
switch# clock set <HH:MM:SS> <DD> <Month in words> <YYYY>
```

For example:

```
switch# clock set 15:58:09 23 September 2002
Mon Sep 23 15:58:09 UTC 2002
```

Where *HH* represents hours in military format (15 for 3 p.m.), *MM* is minutes (58), *SS* is seconds (09), *DD* is the date (23), *Month* is the month in words (September), and *YYYY* is the year (2002).

Note The date and time changes are saved across system resets.

Configuring the Time Zone

You can specify a time zone for the switch.

To specify the local time without the daylight saving time feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

Step 2	`switch(config)# clock timezone <timezone name> <-23 to 23 hours offset from UTC time> <0 to 50 minutes offset from UTC>`
	
	`Example:`
	`switch(config)# clock timezone PST -8 0`
		Sets the time zone with a specified name, specified hours, and specified minutes.
		
		This example sets the time zone to Pacific Standard Time (PST) and offsets the UTC time by negative eight hours and 0 minutes.

Step 3	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 4	`switch# show clock`	Verifies the time zone configuration.
Step 5	`switch# show run`	Displays changes made to the time zone configuration along with other configuration information.

Adjusting for Daylight Saving Time or Summer Time

You can configure your switch to adjust for daylight saving time (or summer time). By default, MDS SAN-OS does not automatically adjust for daylight saving time. You must manually configure the switch to adjust to the daylight saving time.

For example, following U.S. standards, you can have the switch advance the clock one hour at 2:00 a.m. on the first Sunday in April and move back the clock one hour at 2:00 a.m. on the last Sunday in October. You can also explicitly specify the start and end dates and times and whether or not the time adjustment recurs every year.

Note To enable the daylight saving time clock adjustment, follow these steps:In 2007, the U. S. the daylight

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# clock timezone` *timezone_name hour_offset_from_UTC minute_offset_from_UTC* Example: `switch(config)# clock timezone PST -8 0`	Offsets the time zone as specified. This example sets the U.S. Pacific standard offset time as negative 8 hours and 0 minutes.
	`switch(config)# no clock timezone`	Disables the time zone adjustment feature.

Step 3	`switch(config)# clock summer-time` *daylight_timezone_name start_week start_day start_month start_time end_week end_day end_month end_time daylight_offset_inminutes*
	
	Example:
	`switch(config)# clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60`
	`switch(config)#`
		Sets the daylight savings time for a specified time zone.
		The start and end values are as follows:
		• Week ranging from 1 through 5
		• Day ranging from Sunday through Saturday
		• Month ranging from January through December
		The daylight offset ranges from 1 through 1440 minutes, which are added to the start time and deleted time from the end time.
		This example adjusts the daylight savings time for the U.S. Pacific daylight time by 60 minutes starting the second Sunday in March at 2 a.m. and ending the first Sunday in November at 2 a.m.

	`switch(config)# no clock summer-time`	Disables the daylight saving time adjustment feature.
Step 4	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 5	`switch# show running-config | include summer-time`	Verifies the time zone configuration. If *summer-time* is not part of the running configuration, then the switch is not configured for daylight savings time.

saving time adjustment occurs on the second Sunday in March and end on the first Sunday in November. You can update the configuration of your switch to accommodate this change using the following command:

switch(config)# **clock summer-time daylight_timezone_name 2 Sunday March 02:00 1 Sunday November 02:00 60**

Note CFS does not support daylight savings time because a single fabric can span multiple time zones; every switch must be configured individually.

If you want to configure daylight savings time on multiple switches simultaneously, see the RUN CLI command feature in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

NTP Configuration

A Network Time Protocol (NTP) server provides a precise time source (radio clock or atomic clock) to synchronize the system clocks of network devices. NTP is transported over User Datagram Protocol UDP/IP. All NTP communications use Universal Time Coordinated (UTC). An NTP server receives its time from a reference time source, such as a radio clock or atomic clock, attached to the time. NTP distributes this time across the network.

This section includes the following sections:

About NTP

In a large enterprise network, having one time standard for all network devices is critical for management reporting and event logging functions when trying to correlate interacting events logged across multiple devices. Many enterprise customers with extremely mission-critical networks maintain their own stratum-1 NTP source.

Time synchronization happens when several frames are exchanged between clients and servers. The switches in client mode know the address of one or more NTP servers. The servers act as the time source and receive client synchronization requests.

By configuring an IP address as a peer, the switch will obtain and provide time as required. The peer is capable of providing time on its own and is capable of having a server configured. If both these instances point to different time servers, your NTP service is more reliable. Thus, even if the active server link is lost, you can still maintain the right time due to the presence of the peer.

Tip If an active server fails, a configured peer helps in providing the NTP time. Provide a direct NTP server association and configure a peer to ensure backup support if the active server fails.

If you only configure a peer, the most accurate peer takes on the role of the NTP server and the other peer(s) acts as a peer(s). Both machines end at the right time if they have the right time source or if they point to the right NTP source.

NTP Configuration Guidelines

The following guidelines apply to all NTP configurations:

- You should have a peer association with another switch only when you are sure that your clock is reliable (which means that you are a client of a reliable NTP server).

- A peer configured alone takes on the role of a server and should be used as backup. If you have two servers, then you can have several switches point to one server, and the remaining switches to the other server. Then you would configure peer association between these two sets. This forces the clock to be more reliable.
- If you only have one server, it's better for all the switches to have a client association with that server.

Not even a server down time will affect well-configured switches in the network. Figure 5-3 displays a network with two NTP stratum 2 servers and two switches.

Figure 5-3 NTP Peer and Server Association

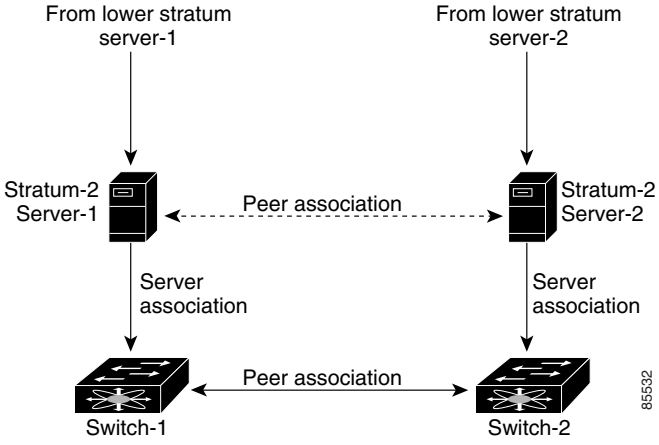

In this configuration, the switches were configured as follows:

- Stratum 2 Server 1
  - IPv4 address–10.10.10.10
  - Stratum–2 Server-2
  - IPv4 address–10.10.10.9
- Switch 1 IPv4 address–10.10.10.1
- Switch 1 NTP configuration
  - NTP server 10.10.10.10
  - NTP peer 10.10.10.2
- Switch 2 IPv4 address–10.10.10.2
- Switch 2 NTP configuration
  - NTP server 10.10.10.9
  - NTP peer 10.10.10.1

Configuring NTP

You can configure NTP using either IPv4 addresses, IPv6 addresses, or DNS names.

To configure NTP in a server association using IPv4 addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp server 10.10.10.10`	Forms a server association with a server.
Step 3	`switch(config)# ntp peer 10.20.10.0`	Forms a peer association with a peer. You can specify multiple associations.
Step 4	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 5	`switch# copy running-config startup-config`	Saves your configuration changes to NVRAM. **Tip** This is one instance where you can save the configuration as a result of an NTP configuration change. You can issue this command at any time.

Step 6	`switch# show ntp peers`
	`----------`
	`Peer IP Address Serv/Peer`
	`----------`
	`10.20.10.0 Peer (configured)`
	`10.10.10.10 Server (configured)`
		Displays the configured server and peer associations.
		Note A domain name is resolved only when you have a DNS server configured.

To configure NTP in a server association using IPv6 addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp server 2001:db8:800:200c::4101`	Forms a server association with a server.
Step 3	`switch(config)# ntp peer 2001:db8:800:200c::417a`	Forms a peer association with a peer. You can specify multiple associations.
Step 4	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 5	`switch# copy running-config startup-config`	Saves your configuration changes to NVRAM. **Tip** This is one instance where you can save the configuration as a result of an NTP configuration change. You can issue this command at any time.

Step 6	`switch# show ntp peers`
	`----------`
	`Peer IP Address Serv/Peer`
	`----------`
	`2001:db8:800:200c::417a Peer (configured)`
	`2001:db8:800:200c::4101 Server (configured)`
		Displays the configured server and peer associations.
		Note A domain name is resolved only when you have a DNS server configured.

To configure NTP in a server association using DNS names, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp server NtpServer`	Forms a server association with a server.

	Command	Purpose
Step 3	`switch(config)# ntp peer NtpPeer`	Forms a peer association with a peer. You can specify multiple associations.
Step 4	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 5	`switch# copy running-config startup-config`	Saves your configuration changes to NVRAM. **Tip** This is one instance where you can save the configuration as a result of an NTP configuration change. You can issue this command at any time.

Step 6	`switch# show ntp peers`
	`----------------------------------------------`
	`  Peer IP Address  Serv/Peer`
	`----------------------------------------------`
	`  NtpPeer  Peer (configured)`
	`  NtpServer  Server (configured)`
		Displays the configured server and peer associations.

NTP CFS Distribution

You can enable NTP fabric distribution for all Cisco MDS switches in the fabric. When you perform NTP configurations, and distribution is enabled, the entire server/peer configuration is distributed to all the switches in the fabric.

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

See to Chapter 6, "Using the CFS Infrastructure," for more information on the CFS application.

This section includes the following sections:

Enabling NTP Distribution

To enable NTP configuration fabric distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp distribute`	Enables NTP configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.
	`switch(config)# no ntp distribute`	Disables (default) NTP configuration distribution to all switches in the fabric.

Committing NTP Configuration Changes

When you commit the NTP 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 NTP configuration changes without implementing the session feature, the NTP configurations are distributed to all the switches in the fabric.

To commit the NTP configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp commit`	Distributes the NTP 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 NTP Configuration Changes

After making the configuration changes, you can choose to discard the changes or to commit them. In either case, the lock is released.

To discard NTP configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ntp abort`	Discards the NTP configuration changes in the pending database and releases the fabric lock.

Releasing Fabric Session Lock

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

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

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

```
switch# clear ntp session
```

Database Merge Guidelines

When merging two fabrics, follow these guidelines:

- Be aware that the merge is a union of the existing and the received database in each switch in the fabric.
- Do not configure an IP address as a server on one switch and as a peer on another switch. The merge can fail if this configuration exists.
- Verify that the union of the databases does not exceed the maximum limit of 64.

See to the "CFS Merge Support" section on page 6-8 for detailed concepts.

NTP Session Status Verification

To verify the status of the NTP session, use the s**how ntp session-status** command.

```
switch# show ntp session-status
last-action : Distribution Enable    Result : Success
```

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface (mgmt0), but first you must configure some IP parameters so that the switch is reachable. You can manually configure the management interface from the CLI. You can configure the mgmt 0 interface with either IPv4 address parameters or an IPv6 address.

On director class switches, a single IP address is used to manage the switch. The active supervisor module's mgmt0 interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.

The management port (mgmt0) is autosensing and operates in full duplex mode at a speed of 10/100/1000 Mbps (1000 Mbps is only available on the Supervisor-2 module). Autosensing supports both the speed and the duplex mode. On a Supervisor-1 module, the default speed is 100 Mbps and the default duplex mode is auto. On a Supervisor-2 module, the default speed is auto and the default duplex mode is auto.

Note Before you begin to configure the management interface manually, obtain the switch's IPv4 address and IPv4 subnet mask or the IPv6 address. Also make sure the console cable is connected to the console port.

Obtaining Remote Management Access

In some cases, a switch interface might be administratively shut down. You can check the status of an interface at any time by using the **show interface mgmt 0** command.

To obtain remote management access using IPv4 addressing parameters, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode. You can also abbreviate the command to **config t**.
Step 2	`switch(config)# interface mgmt 0` `switch(config-if)#`	Enters the interface configuration mode on the specified interface (mgmt0). You can use the console interface on the switch to configure the management Ethernet interface.
Step 3	`switch(config-if)# ip address 1.1.1.0 255.255.255.0`	Specifies the IPv4 address and IPv4 subnet mask.
Step 4	`switch(config-if)# switchport speed 100`	Configures the port speed in Mbps. Valid values are **10**, **100**, and **1000** (Supervisor-2 module only).
Step 5	`switch(config-if)# no shutdown`	Enables the interface.
Step 6	`switch(config-if)# exit`	Returns to configuration mode.
Step 7	`switch(config)# ip default-gateway 1.1.1.1`	Configures the IPv4 default gateway address.

To obtain remote management access using IPv6 addressing parameters, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode. You can also abbreviate the command to **config t**.
Step 2	`switch(config)# interface mgmt 0` `switch(config-if)#`	Enters the interface configuration mode on the specified interface (mgmt0). You can use the console interface on the switch to configure the management Ethernet interface.
Step 3	`switch(config-if)# ipv6 address 2001:0DB8:800:200C::417A/64`	Specifies the IPv6 address and prefix length.
	`switch(config-if)# ipv6 address autoconfig`	Specifies IPv6 autoconfiguration.
Step 4	`switch(config-if)# switchport speed 100`	Configures the port speed in Mbps. Valid values are **10**, **100**, and **1000** (Supervisor-2 module only).
Step 5	`switch(config-if)# no shutdown`	Enables the interface.
Step 6	`switch(config-if)# exit` `switch(config)#`	Returns to configuration mode.

Using the force Option During Shutdown

When you try to shut down a management interface (mgmt0), a follow-up message confirms your action before performing the operation. You can use the **force** option to bypass this confirmation. The following example shuts down the interface without using the **force** option:

```
switch# config t
switch(config)# interface mgmt 0
switch(config-if)# shutdown
Shutting down this interface will drop all telnet sessions.
Do you wish to continue (y/n)? y
```

The following example shuts down the interface using the **force** option:

```
switch# config t
switch(config)# interface mgmt 0
switch(config-if)# shutdown force
```

Note You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Default Gateway Configuration

The supervisor module sends IP packets with unresolved destination IPv4 addresses to the default gateway (see Figure 5-4).

Figure 5-4 Default Gateway

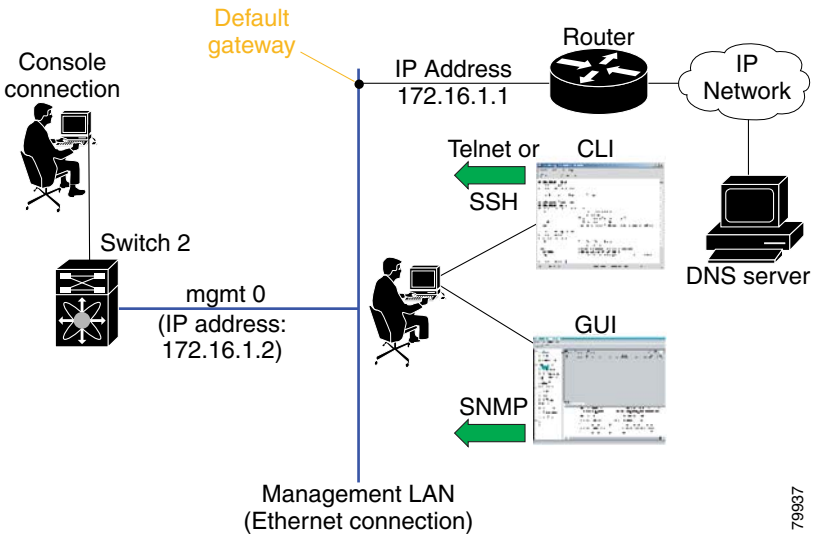

Configuring the Default Gateway

To configure the IPv4 address of the default gateway, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip default-gateway 172.16.1.1`	Configures the 172.16.1.1 IPv4 address.

Telnet Server Connection

The Telnet server is enabled by default on all switches in the Cisco MDS 9000 Family. If you require a secure SSH connection, you need to disable the default Telnet connection and then enable the SSH connection (see the "Generating the SSH Server Key Pair" section on page 37-16).

Note For information on connecting a terminal to the supervisor module console port, refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Tip A maximum of 16 sessions are allowed in any switch in the Cisco MDS 9500 Series or the Cisco MDS 9200 Series.

Make sure the terminal is connected to the switch and that the switch and terminal are both powered on.

Disabling a Telnet Connection

To disable Telnet connections to the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# no telnet server enable` `updated`	Disables the Telnet server.
	`switch(config)# telnet server enable` `updated`	Enables (default) the Telnet server to return a Telnet connection from a secure SSH connection.

Configuring Console Port Settings

The console port is an asynchronous serial port that enables switches in the Cisco MDS 9000 Family to be set up for initial configuration through a standard RS-232 port with an RJ-45 connector. Any device connected to this port must be capable of asynchronous transmission. Connection to a terminal requires a terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.

Caution The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To configure the console port parameters from the console terminal, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line console` `switch(config-console)#`	Enters the line console configuration mode.
Step 3	`switch(config-console)# speed 9600`	Configures the port speed for the serial console. The default console baud rate is 9600 baud. The valid values for Supervisor-1 modules are between 110 and 115200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). The valid values for Supervisor-2 modules are 9600, 19200, 38400, and 115200. Be sure to specify one of these exact values.
Step 4	`switch(config-console)# databits 8`	Configures the data bits for the console connection. The default is 8 data bits and the valid range is between 5 and 8 data bits.
Step 5	`switch(config-console)# stopbits 1`	Configures the stop bits for the console connection. The default is 1 stop bit and the valid values are 1 or 2 stop bits.
Step 6	`switch(config-console)# parity none`	Configures the parity for the console connection. The default is no parity and the valid values are even or odd parity.

Verifying Console Port Settings

Use the **show line console** command to verify the configured console settings. This command also displays problems that may have occurred along with the other registration statistics.

The following example displays output from an MDS switch with a Supervisor-1 module.

```
switch# show line console
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATQ0V1H0S0=1\015
    Statistics:  tx:12842     rx:366     Register Bits:RTS|CTS|DTR|DSR|CD|RI
```

The following example displays output from an MDS switch with a Supervisor-2 module.

```
switch# show line console
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q0V1&D0&C0S0=1\015
    Statistics:  tx:12842     rx:366     Register Bits:RTS|CTS|DTR|DSR|CD|RI
```

Configuring COM1 Port Settings

A COM1 port is an RS-232 port with a DB-9 interface that enables you to connect to an external serial communication device such as a modem. Connection to a terminal requires the terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.

To configure the COM1 port settings, follow these steps:

	Command	Description
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-com1)#`	Enters the COM1 port configuration mode.
Step 3	`switch(config-com1)# speed 9600`	Configures the port speed for the COM1 connection. The default console baud rate is 9600 baud. The valid range is between 110 and 115,200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). Be sure to specify one of these exact values. **Note** This configuration depends on the incoming speed of the modem connected to COM1.
Step 4	`switch(config-com1)# databits 8`	Configures the data bits for the COM1 connection. The default is 8 data bits and the valid range is between 5 and 8 data bits.

	Command	Description
Step 5	`switch(config-com1)# stopbits 1`	Configures the stop bits for the COM1 connection. The default is 1 stop bits and the valid values are 1 or 2 stop bits.
Step 6	`switch(config-com1)# parity none`	Configures the parity for the COM1 connection. The default is no parity and the valid values are even or odd parity.
Step 7	`switch(config-com1)# no flowcontrol hardware`	Disables hardware flow control. By default, hardware flow control is enabled on all switches in the Cisco 9000 Family. When enabled, this option is useful in protecting data loss at higher baud rates. **Note** This option is only available through the COM1 port.

Verifying COM1 Port Settings

Use the **show line com1** command to verify the configured COM1 settings. This command also displays problems that may have occurred along with the other registration statistics.

The following example displays output from an MDS switch with a Supervisor-1 module.

```
switch# show line com1
line Aux:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATQ0V1H0S0=1\015
    Statistics:  tx:17     rx:0     Register Bits:RTS|DTR
```

The following example displays output from an MDS switch with a Supervisor-2 module.

```
switch# show line com1
line Aux:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q0V1&D0&C0S0=1\015
    Statistics:  tx:17     rx:0     Register Bits:RTS|DTR
```

Configuring Modem Connections

Modems can only be configured if you are connected to the console or COM1 ports. A modem connection to a switch in the Cisco MDS 9000 Family does not affect switch functionality.

Note If you plan on connecting a modem to the console port or the COM1 port of a switch in the Cisco MDS 9000 Family, refer to the *Cisco MDS 9500 Series Hardware Installation Guide* or the *Cisco MDS 9200 Series Hardware Installation Guide*. COM1 ports are not available on switches in the Cisco MDS 9100 Series. Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Guidelines to Configure Modems

Tip We recommend you use the COM1 port to connect the modem from any director in the Cisco MDS 9500 Series or any switch in the Cisco MDS 9200 Series.

The following guidelines apply to modem configurations:

- The following modems were tested to work in the Cisco SAN-OS environment using Supervisor-1 modules:
  - MultiTech MT2834BA (http://www.multitech.com/PRODUCTS/Families/MultiModemII/)
  - Hayes Accura V.92 (http://www.hayesmicro.com/Products/accura-prod-v92.htm)
- The following modems were tested to work in the Cisco SAN-OS environment using Supervisor-2 modules:
  - Hayes Accura V.92 (http://www.hayesmicro.com/Products/accura-prod-v92.htm)
  - Zoom/FaxModem 56K Dualmode Model 2949 (http://www.zoom.com/products/dial_up_external_serial.html)
  - Multitech MT2834 BA 33.6K (http://www.multitech.com/PRODUCTS/Families/CC1600-Series/)

    Note On the Multitech MT2834 BA 33.6K set the DIP switch1 (pin1), also known as the DTR-pin, to the DOWN position to enable the DTR signal (or set it to ON). You must connect the modem before attempting to configure it.

  - USRobotics Model 5686 V.92 (http://www.usr.com/products/home/home-product.asp?sku=USR5686E)

    Note On the USRobotics Model 5686 V.92 set the DIP switch1 (pin1), also known as the DTR-pin, to the DOWN position to enable the DTR signal (or set it to ON). You must connect the modem before attempting to configure it.

- Do not connect a modem to the console port while the system is booting.

  Follow the procedure specified in the "Initializing a Modem in a Powered-On Switch" section on page 5-34.

Enabling Modem Connections

To configure a modem connection through the COM1 port, follow these steps:

	Command	Command
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-com1)#`	Enters the COM1 port configuration mode.
Step 3	`switch(config-com1)# modem in`	Enables the COM1 port to only connect to a modem.
	`switch(config-com1)# no modem in`	Disables (default) the current modem from executing its functions.

To configure a modem connection through the console port, follow these steps:

	Command	Command
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line console` `switch(config-console)#`	Enters the console port configuration mode.
Step 3	`switch(config-console)# modem in`	Enables the console port to only connect to a modem.
	`switch(config-console)# no modem in`	Disables (default) the current modem from executing its functions.

Configuring the Initialization String

Switches in the Cisco MDS 9500 Series and the Cisco MDS 9200 Series have a default initialization string (`ATE0Q1&D2&C1S0=1\015`) to detect connected modems. The default string detects connected modems supported by Cisco Systems. The default string contents for Supervisor-1 modules are as follows:

- AT—Attention
- E0 (required)—No echo
- Q1—Result code on
- &D2—Normal data terminal ready (DTR) option
- &C1—Enable tracking the state of the data carrier
- S0=1—Pick up after one ring
- \015 (required)—Carriage return in octal

The default string contents for Supervisor-2 modules are as follows:

- AT—Attention
- E0 (required)—No echo
- Q0—Result code on
- V1—Display result codes as text
- &D0—Data terminal ready (DTR) on

- &C0—Data carrier detect (DCD) on
- S0=1—Pick up after one ring

You may retain the default string or change it to another string (80 character limit) using the **user-input** option. This option is provided if you prefer to use a modem that is not supported or tested by Cisco systems. If you change the string, the changes you make are permanent and remain in effect unless you change them again. Rebooting the system or restarting the CLI does not change the modem initialization string. The switch is not affected even if the modem is not functioning.

Tip We recommend you use the default initialization string. If the required options are not provided in the user-input string, the initialization string is not processed.

The modem initialization string usage depends on the modem state when the switch boots:

- If the modem is already attached to the switch during boot-up, the default initialization string is written to the modem (see the "Configuring the Default Initialization String" section on page 5-33).
- If the modem is not attached to the switch during boot-up, then attach the modem as outlined in the Cisco MDS 9000 Family Hardware Installation Guide (depending on the product), and follow the procedure provided in this section (see the "Configuring a User-Specified Initialization String" section on page 5-34).

Note You can perform the configuration specified in this section only if you are connected to the console port or the COM1 port.

Configuring the Default Initialization String

To configure the default initialization string through the COM1 port, follow these steps:

	Command	Command
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-com1)#`	Enters the COM1 port configuration mode.
Step 3	`switch(config-com1)# modem init-string default`	Writes the default initialization string to the modem.

To configure the default initialization string through the console port, follow these steps:

	Command	Command
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-console)#`	Enters the console port configuration mode.
Step 3	`switch(config-console)# modem init-string default`	Writes the default initialization string to the modem.

Configuring a User-Specified Initialization String

To configure a user-specified initialization string through the COM1 port, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-com1)#`	Enters the COM1 port configuration mode.
Step 3	`switch(config-com1)# modem set-string user-input ATE0Q1&D2&C1S0=3\015`	Assigns the user-specified initialization string for a Supervisor-1 module to its corresponding profile. **Note** You must first set the user-input string before initializing the string.
	`switch(config-com1)# modem set-string user-input ATE0Q0V1&D0&C0S0=1`	Assigns the user-specified initialization string for a Supervisor-2 module to its corresponding profile.
	`switch(config-com1)# no modem set-string`	Reverts the configured initialization string to the factory default string.
Step 4	`switch(config-com1)# modem init-string user-input`	Writes the user-specified initialization string to the modem.

To configure a user-specified initialization string through the console port, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# line com1` `switch(config-console)#`	Enters the console port configuration mode.
Step 3	`switch(config-console)# modem set-string user-input ATE0Q1&D2&C1S0=3\015`	Assigns the user-specified initialization string to its corresponding profile. **Note** You must first set the user-input string before initializing the string.
	`switch(config-com1)# modem set-string user-input ATE0Q0V1&D0&C0S0=1`	Assigns the user-specified initialization string for a Supervisor-2 module to its corresponding profile.
	`switch(config-com1)# no modem set-string`	Reverts the configured initialization string to the factory default string.
Step 4	`switch(config-console)# modem init-string user-input`	Writes the user-specified initialization string to the modem.

Initializing a Modem in a Powered-On Switch

When a switch is already powered-on and the modem is later connected to either the console port or the COM1 port, you can initialize the modem using the **modem connect line** command in EXEC mode. You can specify the **com1** option if the modem is connected to the COM1 port, or the **console** option if the modem is connected to the console.

To connect a modem to a switch that is already powered on, follow these steps.

Step 1 Wait until the system has completed the boot sequence and the system image is running.

Step 2 Connect the modem to the switch as specified in the *Cisco MDS 9500 Series Hardware Guide* or the *Cisco MDS 9200 Series Hardware Installation Guide*.

Step 3 Initialize the modem using the **modem connect line** command in EXEC mode.

Verifying the Modem Connection Configuration

Use the **show line** command to verify the configured modem settings.

The following example displays output from an MDS switch with a Supervisor-1 module.

```
switch# show line
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q1&D2&C1S0=1\015
    Statistics:  tx:12842     rx:366     Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q1&D2&C1S0=1\015
    Statistics:  tx:17     rx:0     Register Bits:RTS|DTR
```

The following example displays output from an MDS switch with a Supervisor-2 module.

```
switch# show line
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q0V1&D0&C0S0=1
    Statistics:  tx:12842     rx:366     Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Enable
    Modem Init-String -
        default : ATE0Q0V1&D0&C0S0=1
    Statistics:  tx:17     rx:0     Register Bits:RTS|DTR
```

Configuring CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of Layer 3 protocols. Cisco devices that receive the CDP packets cache the information to make it is accessible through the CLI and SNMP.

CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS and MPS-14/2 modules. The CDP daemon is restartable and switchable. The running and startup configurations are available across restarts and switchovers.

CDP version 1 (v1) and version 2 (v2) are supported in Cisco MDS 9000 Family switches. CDP packets with any other version number are silently discarded when received.

When the interface link is established, CDP is enabled by default and three CDP packets are sent at one-second intervals. Following this, the CDP frames are sent at the globally configured refresh interval.

To globally disable the CDP, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# no cdp enable`
	`Operation in progress. Please check global parameters`
	`switch(config-console)#`
		Disables the CDP protocol on the switch. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices.

	`switch(config)# cdp enable`
	`Operation in progress. Please check global parameters`
	`switch(config)#`
		Enables (default) the CDP protocol on the switch. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time.

To disable the CDP protocol on a specific interface, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigbitethernet 3/8` `switch(config-if)#`	Configures the Gigabit Ethernet interface for the module in slot 3 port 8.

Step 3	`switch(config-if)# no cdp enable`
	`Operation in progress. Please check interface parameters`
	`switch(config-console)#`
		Disables the CDP protocol on the selected interface. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices.

	`switch(config-if)# cdp enable`
	`Operation in progress. Please check interface parameters`
	`switch(config)#`
		Enables (default) the CDP protocol on the selected interface. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time.

To globally configure the refresh time interval for the CDP protocol, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cdp timer 100` `switch(config)#`	Sets the refresh time interval in seconds. The default is 60 seconds and the valid range is from 5 to 255 seconds.
	`switch(config)# no cdp timer 100` `switch(config)#`	Reverts the refresh time interval to the factory default of 60 seconds.

To globally configure the hold time advertised in CDP packets, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cdp holdtime 200`	Sets the hold time advertised in CDP packets in seconds. The default is 180 seconds and the valid range is from 10 to 255 seconds.
	`switch(config)# no cdp holdtime 200`	Reverts the hold time to the factory default of 180 seconds.

To globally configure the CDP version, follow these steps:

	Command	Command
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cdp advertise v1`	Sets the CDP version to be used. The default is version 2 (v2). The valid options are **v1** and **v2**.
	`switch(config)# no advertise v1`	Reverts the version to the factory default of **v2**.

Clearing CDP Counters and Tables

Use the **clear cdp counters** command to clear CDP traffic counters for all interfaces. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces).

```
switch# clear cdp counters
```

Use the **clear cdp table** command to clear neighboring CDP entries for all interfaces. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces).

```
switch# clear cdp table interface gigabitethernet 4/1
```

Displaying CDP Information

Use the **show cdp** command to display CDP entries. See Examples 5-1 to 5-11.

Example 5-1 Displays All CDP Capable Interfaces and Parameters

```
switch# show cdp all
GigabitEthernet4/1 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
GigabitEthernet4/8 is down
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
mgmt0 is up
    CDP enabled on interface
    Sending CDP packets every 100 seconds
    Holdtime is 200 seconds
```

Example 5-2 Displays All CDP Neighbor Entries

```
switch# show cdp entry all
----------------------------------------
Device ID:069038747(Kiowa3)
Entry address(es):
    IP Address: 172.22.92.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 5/22
Holdtime: 136 sec

Version:
WS-C5500 Software, Version McpSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems

Advertisement Version: 1
```

Example 5-3 Displays the Specified CDP Neighbor

```
switch# show cdp entry name 0
----------------------------------------
Device ID:0
Entry address(es):
    IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
```

Example 5-4 Displays Global CDP Parameters

```
switch# show cdp global
Global CDP information:
    CDP enabled globally
    Sending CDP packets every 60 seconds
```

```
    Sending a holdtime value of 180 seconds
    Sending CDPv2 advertisements is enabled
```

Example 5-5 Displays CDP Parameters for the Management Interface

```
switch# show cdp interface mgmt 0
mgmt0 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
```

Example 5-6 Displays CDP Parameters for the Gigabit Ethernet Interface

```
switch# show cdp interface gigabitethernet 4/1
GigabitEthernet4/1 is up
    CDP enabled on interface
    Sending CDP packets every 80 seconds
    Holdtime is 200 seconds
```

Example 5-7 Displays CDP Neighbors (in brief)

```
switch# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce   Hldtme  Capability  Platform      Port ID
0                Gig4/1          135     H           DS-X9530-SF1- Gig4/1
069038732(Kiowa2 mgmt0           132     T S         WS-C5500      3/3/11
069038747(Kiowa3 mgmt0           156     T S         WS-C5500      6/20
069038747(Kiowa3 mgmt0           158     T S         WS-C5500      5/22
```

Example 5-8 Displays CDP Neighbors (in detail)

```
switch# show CDP neighbor detail
----------------------------------------
Device ID:0
Entry address(es):
    IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 162 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
----------------------------------------
Device ID:069038732(Kiowa2)
Entry address(es):
    IP Address: 172.22.91.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 3/11
Holdtime: 132 sec

Version:
WS-C5500 Software, Version McpSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems
Advertisement Version: 1
```

Example 5-9 Displays the Specified CDP Neighbor (in detail)

```
switch# show CDP neighbors interface gigabitethernet 4/1 detail
----------------------------------------
Device ID:0
Entry address(es):
    IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
```

Example 5-10 Displays CDP Traffic Statistics for the Management Interface

```
switch# show cdp traffic interface mgmt 0
----------------------------------------
Traffic statistics for mgmt0
Input Statistics:
    Total Packets: 1148
    Valid CDP Packets: 1148
        CDP v1 Packets: 1148
        CDP v2 Packets: 0
    Invalid CDP Packets: 0
        Unsupported Version: 0
        Checksum Errors: 0
        Malformed Packets: 0
Output Statistics:
    Total Packets: 2329
        CDP v1 Packets: 1164
        CDP v2 Packets: 1165
    Send Errors: 0
```

Example 5-11 Displays CDP Traffic Statistics for the Gigabit Ethernet Interface

```
switch# show cdp traffic interface gigabitethernet 4/1
----------------------------------------
Traffic statistics for GigabitEthernet4/1
Input Statistics:
    Total Packets: 674
    Valid CDP Packets: 674
        CDP v1 Packets: 0
        CDP v2 Packets: 674
    Invalid CDP Packets: 0
        Unsupported Version: 0
        Checksum Errors: 0
        Malformed Packets: 0

Output Statistics:
    Total Packets: 674
        CDP v1 Packets: 0
        CDP v2 Packets: 674
    Send Errors: 0
```

CHAPTER 6

Using the CFS Infrastructure

The Cisco MDS SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

Several Cisco MDS SAN-OS applications use the CFS infrastructure to maintain and distribute the contents of a particular application's database.

This chapter contains the following sections:

About CFS

Many features in the Cisco MDS switches require configuration synchronization in all switches in the fabric. Maintaining configuration synchronization across a fabric is important to maintain fabric consistency. In the absence of a common infrastructure, such synchronization is achieved through manual configuration at each switch in the fabric. This process is tedious and error prone.

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the applications. CFS has the ability to discover CFS capable switches in the fabric and discovering application capabilities in all CFS capable switches.

This section includes the following topics:

Cisco SAN-OS Features Using CFS

The following Cisco SAN-OS features use the CFS infrastructure:

- NTP (see the "NTP CFS Distribution" section on page 5-23).
- Dynamic Port VSAN Membership (see the "DPVM Database Distribution" section on page 21-5).
- Distributed Device Alias Services (see the "Device Alias Databases" section on page 24-3).
- IVR topology (see the "Database Merge Guidelines" section on page 22-37).
- SAN device virtualization (see the "Configuring SDV" section on page 20-4).
- TACACS+ and RADIUS (see the "AAA Server Distribution" section on page 33-30).
- User and administrator roles (see the "Role-Based Authorization" section on page 37-1).
- Port security (see the "Port Security Configuration Distribution" section on page 38-11).
- iSNS (see the "iSNS" section on page 42-82).
- Call Home (see the "Call Home Configuration Distribution" section on page 54-13).
- Syslog (see the "System Message Logging Configuration Distribution" section on page 53-8).
- fctimer (see the "About fctimer Distribution" section on page 29-4).
- SCSI flow services (see the "Configuring SCSI Flow Services" section on page 47-3).
- Saving startup configurations in the fabric using the Fabric Startup Configuration Manager (FSCM) (see the "Saving Startup Configurations in the Fabric" section on page 8-4).
- Allowed domain ID lists (see the"About Allowed Domain ID Lists" section on page 17-10).
- RSCN timer (see the "Configuring the RSCN Timer" section on page 26-10).
- iSLB (see the "About iSLB Configuration Distribution Using CFS" section on page 42-57).

CFS Features

CFS has the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- Three scopes of distribution.
  - Logical scope: The distribution occurs within the scope of a VSAN.
  - Physical scope: The distribution spans the entire physical topology.
  - Over a selected set of VSANs: Some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.

- Three modes of distribution.
  - Coordinated distributions: Only one distribution is allowed in the fabric at any given time.
  - Uncoordinated distributions: Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
  - Unrestricted uncoordinated distributions: Multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Supports a merge protocol that facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the FC2 layer and is peer-to-peer with not client-server relationship. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

CFS can also use IP to send information to other switches (see the "CFS Distribution over IP" section on page 6-11).

Applications that use CFS are completely unaware of the lower layer transport.

CFS Distribution Scopes

Different applications on the Cisco MDS 9000 Family switches need to distribute the configuration at various levels:

- VSAN level (logical scope)

  Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.
- Physical topology level (physical scope)

  Applications might need to distribute the configuration to the entire physical topology spanning several VSANs. Such applications include NTP and DPVM (WWN based VSAN), which are independent of VSANs.
- Betweenselected switches

  Applications might only operate between selected switches in the fabric. An example application is SCSI Flow Services, which operates between two switches.

CFS Distribution Modes

CFS supports different distribution modes to support different application requirements: coordinated and uncoordinated distributions. Both modes are mutually exclusive. Only one mode is allowed at any given time.

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. An example is local device registrations such as iSNS. Parallel uncoordinated distributions are allowed for an application.

Coordinated Distribution

Coordinated distributions can have only one application distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the application anywhere in the fabric. A coordinated distribution consists of three stages:

1. A fabric lock is acquired.
2. The configuration is distributed and committed.
3. The fabric lock is released.

Coordinated distribution has two variants:

- CFS driven —The stages are executed by CFS in response to an application request without intervention from the application.
- Application driven—The stages are under the complete control of the application.

  Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Unrestricted Uncoordinated Distributions

Unrestricted uncoordinated distributions allow multiple parallel distributions in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

Disabling CFS Distribution on a Switch

By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation.

You can globally disable CFS on a switch, including CFS over IP, to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

To globally disable or enable CFS distribution on a switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no cfs distribute`	Globally disables CFS distribution for all applications on the switch, including CFS over IP.
	`switch(config)# cfs distribute`	Enables (default) CFS distribution on the switch.

Verifying CFS Distribution Status

The **show cfs status** command displays the status of CFS distribution on the switch.

```
switch# show cfs status
Fabric distribution Enabled
```

CFS Application Requirements

All switches in the fabric must be CFS capable. A Cisco MDS 9000 Family switch is CFS capable if it is running Cisco SAN-OS Release 2.0(1b) or later. Switches that are not CFS capable do not receive distributions and result in part of the fabric not receiving the intended distribution.

CFS has the following requirements:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the fabric.
- Pending database—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the pending database overwrites the configuration database (also known as the active database or the effective database).
- CFS distribution enabled or disabled on a per-application basis—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the fabric.
- Explicit CFS commit—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database, to distribute the new database to the fabric, and to release the fabric lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.

Enabling CFS for an Application

All CFS based applications provide an option to enable or disable the distribution capabilities. Features that existed prior to Cisco SAN-OS Release 2.0(1b) have the distribution capability disabled by default and must have distribution capabilities enabled explicitly.

Applications introduced in Cisco SAN-OS Release 2.0(1b) or later have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

Verifying Application Registration Status

The **show cfs application** command displays the applications that are currently registered with CFS. The first column displays the application name. The second column indicates whether the application is enabled or disabled for distribution (`enabled` or `disabled`). The last column indicates the scope of distribution for the application (`logical`, `physical`, or `both`).

Note The **show cfs application** command only displays applications registered with CFS. Conditional services that use CFS do not appear in the output unless these services are running.

```
switch# show cfs application
----------------------------------------------
 Application    Enabled   Scope
----------------------------------------------
 ntp            No        Physical-all
 fscm           Yes       Physical-fc
 islb           No        Physical-fc
 role           No        Physical-all
 rscn           No        Logical
 radius         No        Physical-all
 fctimer        No        Physical-fc
 syslogd        No        Physical-all
 callhome       No        Physical-all
 fcdomain       No        Logical
 device-alias   Yes       Physical-fc

Total number of entries = 11
```

The **show cfs application name** command displays the details for a particular application. It displays the enabled/disabled state, timeout as registered with CFS, merge capability (if it has registered with CFS for merge support), and lastly the distribution scope.

```
switch# show cfs application name ntp

 Enabled        : Yes
 Timeout        : 5s
 Merge Capable  : Yes
 Scope          : Physical
```

Locking the Fabric

When you configure (first time configuration) a Cisco SAN-OS feature (or application) that uses the CFS infrastructure, that feature starts a CFS session and locks the fabric. When a fabric is locked, the Cisco SAN-OS software does not allow any configuration changes from a switch, other than the switch holding the lock, to this Cisco SAN-OS feature and issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a fabric lock but forget to end the session, an administrator can clear the session. If you lock a fabric at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

Verifying CFS Lock Status

The **show cfs lock** command displays all the locks that are currently acquired by any application. For each application the command displays the application name and scope of the lock taken. If the application lock is taken in the physical scope, then this command displays the switch WWN, IP address, user name, and user type of the lock holder. If the application is taken in the logical scope, then this command displays the VSAN in which the lock is taken, the domain, IP address, user name, and user type of the lock holder.

```
switch# show cfs lock

Application: ntp
Scope      : Physical
--------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
--------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3
 Total number of entries = 1

Application: port-security
Scope      : Logical
----------------------------------------------------------
 VSAN   Domain   IP Address       User Name    User Type
----------------------------------------------------------
 1      238      10.76.100.167    admin        CLI/SNMP v3
 2      211      10.76.100.167    admin        CLI/SNMP v3
 Total number of entries = 2
```

The **show cfs lock name** command displays the lock details similar for the specified application.

Example 6-1 Displays the Lock Information for the Specified Application

```
switch# show cfs lock name ntp
Scope      : Physical
--------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
--------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3

 Total number of entries = 1
```

Committing Changes

A commit operation saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session—only a lock function starts a session. However, an empty commit is allowed if configuration changes are not previously made. In this case, a commit operation results in a session that acquires locks and distributes the current database.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the fabric lock.
- None of the external switches report a successful state—The application considers this state a failure and does not apply the changes to any switch in the fabric. The fabric lock is not released.

You can commit changes for a specified feature byissuing the **commit** command for that feature.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the fabric. Both the abort and commit functions are only supported from the switch from which the fabric lock is acquired.

You can discard changes for a specified feature by using the **abort** command for that feature.

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrite the configuration in the effective database when you commit the changes.

Caution If you do not commit the changes, they are not saved to the running configuration.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information on this MIB.

Clearing a Locked Session

You can clear locks held by an application from any switch in the fabric. This option is provided to rescue you from situations where locks are acquired and not released. This function requires Admin permissions.

Caution Exercise caution when using this function to clear locks in the fabric. Any pending configurations in any switch in the fabric is flushed and lost.

CFS Merge Support

An application keeps the configuration synchronized in a fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification

each time an application peer comes online. If a fabric with M application peers merges with another fabric with N application peers and if an application triggers a merge action on every such notification, a link-up event results in M*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not play any role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

Verifying CFS Merge Status

The **show cfs merge status name** command displays the merge status for a given application. The following example displays the output for an application distributing in logical scope. It shows the merge status in all valid VSANs on the switch. The command output shows the merge status as one of the following: `Success, waiting,` or `Failure` or `In Progress`. In case of a successful merge, all the switches in the fabric are shown under the local fabric. In case of a merge failure or a merge in progress, the local fabric and the remote fabric involved in the merge are indicated separately. The application server in each fabric that is mainly responsible for the merge is indicated by the term `Merge Master`.

```
switch# show cfs merge status name port-security

Logical [VSAN 1] Merge Status: Failed
Local Fabric
-----------------------------------------------------------------
 Domain Switch WWN               IP Address
-----------------------------------------------------------------
 238    20:00:00:05:30:00:6b:9e  10.76.100.167   [Merge Master]

Remote Fabric
-----------------------------------------------------------------
 Domain Switch WWN               IP Address
-----------------------------------------------------------------
 236    20:00:00:0e:d7:00:3c:9e  10.76.100.169    [Merge Master]

Logical [VSAN 2] Merge Status: Success
Local Fabric
-----------------------------------------------------------------
 Domain Switch WWN               IP Address
-----------------------------------------------------------------
 211    20:00:00:05:30:00:6b:9e  10.76.100.167    [Merge Master]
 1      20:00:00:0e:d7:00:3c:9e  10.76.100.169

Logical [VSAN 3] Merge Status: Success
Local Fabric
-----------------------------------------------------------------
 Domain Switch WWN               IP Address
-----------------------------------------------------------------
 221    20:00:00:05:30:00:6b:9e  10.76.100.167    [Merge Master]
 103    20:00:00:0e:d7:00:3c:9e  10.76.100.169
```

The following example of the **show cfs merge status name** command output displays an application using the physical scope with a merge failure. The command uses the specified application name to display the merge status based on the application scope.

```
switch# show cfs merge status name ntp

Physical  Merge Status: Failed
Local Fabric
---------------------------------------------------------
 Switch WWN               IP Address
---------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    [Merge Master]

Remote Fabric
---------------------------------------------------------
 Switch WWN               IP Address
---------------------------------------------------------
 20:00:00:0e:d7:00:3c:9e  10.76.100.169    [Merge Master]
```

The **show cfs peers** command output displays all the switches in the physical fabric in terms of the switch WWN and the IP address. The local switch is indicated as `Local`.

```
switch# show cfs peers

Physical Fabric
-------------------------------------------------
 Switch WWN               IP Address
-------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167   [Local]
 20:00:00:0e:d7:00:3c:9e  10.76.100.169

 Total number of entries = 2
```

The **show cfs peers name** command displays all the peers for which a particular application is registered with CFS. The command output shows all the peers for the physical scope or for each of the valid VSANs on the switch, depending on the application scope. For physical scope, the switch WWNs for all the peers are indicated. The local switch is indicated as `Local`.

```
switch# show cfs peers name ntp

Scope      : Physical
-------------------------------------------------
Switch WWN               IP Address
-------------------------------------------------
20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
20:00:00:05:30:01:1b:c2  172.22.92.215
```

The following example **show cfs peers name** command output displays all the application peers (all switches in which that application is registered). The local switch is indicated as `Local`.

```
switch# show cfs peers name port-security
Scope      : Logical [VSAN 1]
-----------------------------------------------------------
Domain   Switch WWN               IP Address
-----------------------------------------------------------
124      20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
98       20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2


Scope      : Logical [VSAN 3]
```

```
-----------------------------------------------------------
Domain    Switch WWN               IP Address
-----------------------------------------------------------
224       20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
151       20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2
```

CFS Distribution over IP

You can configure CFS to distribute information over IP for networks containing switches that are not reachable over Fibre Channel. CFS distribution over IP supports the following features:

- Physical distribution over an entirely IP network.
- Physical distribution over a hybrid Fibre Channel and IP network with the distribution reaching all switches that are reachable over either Fibre Channel or IP

Note The switch attempts to distribute information over Fibre Channel first and then over the IP network if the first attempt over Fibre Channel fails. CFS does not send duplicate messages if distribution over both IP and Fibre Channel is enabled.

- Distribution over IP version 4 (IPv4) or IP version 6 (IPv6).

Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

- Keep-alive mechanism to detect network topology changes using a configurable multicast address.
- Compatibility with Cisco MDS SAN-OS Release 2.x.
- Distribution for logical scope applications is not supported because the VSAN implementation is limited to Fibre Channel.

Figure 6-1 shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

Figure 6-1 Network Example 1 with Fibre Channel and IP Connections

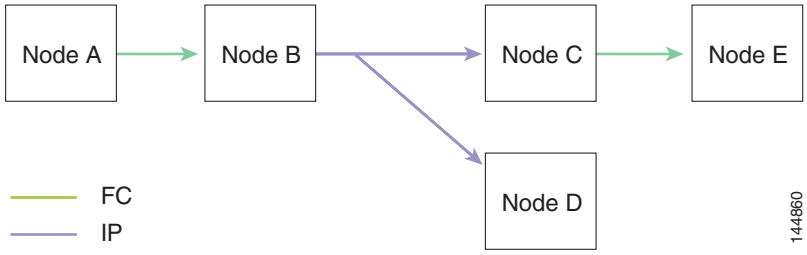

Figure 6-2 is the same as Figure 6-1 except that node C and node D are connected using Fibre Channel. All processes is the same in this example because node B has node C and node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 6-2 Network Example 2 with Fibre Channel and IP Connections

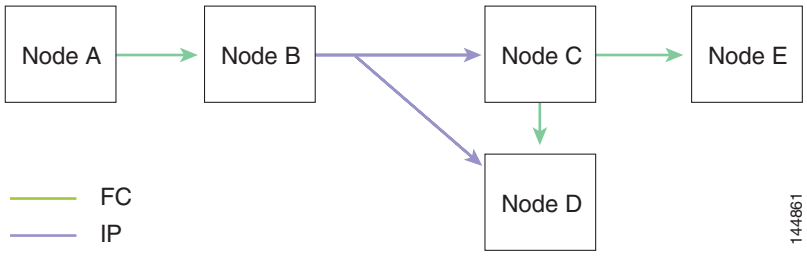

Figure 6-3 is the same as Figure 6-2 except that node D and node E are connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 6-3 Network Example 3 with Fibre Channel and IP Connections

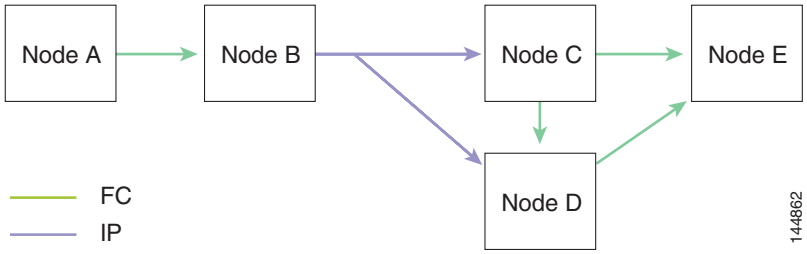

Enabling CFS Over IP

To enable or disable CFS over IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cfs ipv4 distribute`	Globally enables CFS over IPv4 for all applications on the switch.

	`switch(config)# no cfs ipv4 distribute`
	`This will prevent CFS from distributing over IPv4 network.`
	`Are you sure? (y/n) [n] y`
		Disables (default) CFS over IPv4 on the switch.

To enable or disable CFS over IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# cfs ipv6 distribute`	Globally enables CFS over IPv6 for all applications on the switch.
	`switch(config)# no cfs ipv6 distribute`	Disables (default) CFS over IPv6 on the switch.

Verifying the CFS Over IP Configuration

To verify the CFS over IP configuration, use the **show cfs status** command.

```
switch# show cfs status
Fabric distribution Enabled
IP distribution Enabled mode ipv4
```

Configuring IP Multicast Address for CFS over IP

All CFS over IP enabled switches with similar multicast addresses form one CFS over IP fabric. CFS protocol specific distributions, such as the keep-alive mechanism for detecting network topology changes, use the IP multicast address to send and receive information.

Note CFS distributions for application data use directed unicast.

You can configure a CFS over IP multicast address value for either IPv4 or IPv6. The default IPv4 multicast address is 239.255.70.83 and the default IPv6 multicast address is ff13:7743:4653.

To configure an IP multicast address for CFS over IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# cfs ipv4 mcast-address 239.255.1.1`
	`Distribution over this IP type will be affected`
	`Change multicast address for CFS-IP ?`
	`Are you sure? (y/n) [n] y`
		Configures the IPv4 multicast address for CFS distribution over IPv4. The ranges of valid IPv4 addresses are 239.255.0.0 through 239.255.255.255 and 239.192/16 through 239.251/16.

	`switch(config)# no cfs ipv4 mcast-address 239.255.1.1`
	`Distribution over this IP type will be affected`
	`Change multicast address for CFS-IP ?`
	`Are you sure? (y/n) [n] y`
		Reverts to the default IPv4 multicast address for CFS distribution over IPv4. The default IPv4 multicast address for CFS is 239.255.70.83.

To configure an IP multicast address for CFS over IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# cfs ipv6 mcast-address ff15::e244:4754`
	`Distribution over this IP type will be affected`
	`Change multicast address for CFS-IP ?`
	`Are you sure? (y/n) [n] y`
		Configures the IPv6 multicast address for CFS distribution over IPv6. The range of valid IPv6 addresses is ff15::/16 (ff15::0000:0000 through ff15::ffff:ffff) and ff18::/16 (ff18::0000:0000 through ff18::ffff:ffff).

	`switch(config)# no cfs ipv6 mcast-address ff15::e244:4754`
	`Distribution over this IP type will be affected`
	`Change multicast address for CFS-IP ?`
	`Are you sure? (y/n) [n] y`
		Reverts to the default IPv6 multicast address for CFS distribution over IPv6. The default IPv6 multicast address for CFS over IP is ff15::efff:4653.

Verifying IP Multicast Address Configuration for CFS over IP

To verify the IP multicast address configuration for CFS over IP, use the **show cfs status** command.

```
switch# show cfs status
Fabric distribution Enabled
IP distribution Enabled mode ipv4
IPv4 multicast address : 10.1.10.100
IPv6 multicast address : ff13::e244:4754
```

CFS Regions

This section contains the following topics:

About CFS Regions

A CFS region is a user-defined subset of switches for a given feature or application in its physical distribution scope.When a SAN is spanned across a vast geography, you may need to localize or restrict the distribution of certain profiles among a set of switches based on their physical proximity. Before release 3.2.(1) the distribution scope of an application within a SAN was spanned across the entire physical fabric without the ability to confine or limit the distribution to a required set of switches in the fabric. CFS regions enables you to overcome this limitation by allowing you to create CFS regions, that is, multiple islands of distribution within the fabric, for a given CFS feature or application. CFS regions are designed to restrict the distribution of a feature's configuration to a specific set or grouping of switches in a fabric.

Note You can only configure a CFS region on physical switches in a SAN. You cannot configure a CFS region in a VSAN.

Example Scenario: The callhome is an application that triggers alerts to Network Administrators when a situation arises or something abnormal occurs. When the fabric covers many geographies and with multiple Network Administrators who are each responsible for a subset of switches in the fabric, the callhome application sends alerts to all Network Administrators regardless of their location. For the callhome application to send message alerts selectively to Network Administrators, the physical scope of the application has to be fine tuned or narrowed down, which is achieved by implementing CFS regions.

CFS regions are identified by numbers ranging from 0 through 200. Region 0 is reserved as the default region, and contains every switch in the fabric. You can configure regions from 1 through 200. The default region maintains backward compatibility. If there are switches on the same fabric running releases of SAN-OS before release 3.2(1), only features in Region 0 are supported when those switches are synchronized. Features from other regions are ignored when those switches are synchronized.

If the feature is moved, that is, assigned to a new region, its scope is restricted to that region; it ignores all other regions for distribution or merging purposes. The assignment of the region to a feature has precedence in distribution over its initial physical scope.

You can configure a CFS region to distribute configurations for multiple features. However, on a given switch, you can configure only one CFS region at a time to distribute the configuration for a given feature. Once you assign a feature to a CFS region, its configuration cannot be distributed within another CFS region.

Managing CFS Regions

This section describes how to manage a CFS region. A set of commands are used to complete the following tasks:

Creating CFS Regions

To create a CFS region, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cfs region 4`	Creates a region, for example, number 4.

Assigning Applications to CFS Regions

To assign an application on a switch to a region, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cfs region 4`	Creates a region, for example, number 4.
Step 3	`switch(config-cfs-region)# ntp` `switch(config-cfs-region)# callhome`	Adds application(s).

Moving an Application to a Different CFS Region

To move an application for example, from Region 1 (originating region) with ntp and callhome applications assigned to it, to Region 2 (target region), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cfs region 2`	Enters the Region 2.
Step 3	`switch(config-cfs-region)# ntp` `switch(config-cfs-region)# callhome`	Indicates application(s) to be moved into Region 2 that originally belong to Region 1. For example, here, the ntp and callhome applications are moved to Region 2.

Note If you try adding an application to the same region more than once, you see the error message, "Application already present in the same region."

Removing an Application from a Region

Removing an application from a region is the same as moving the application back to the default region or to Region 0, that is, bringing the entire fabric into the scope of distribution for the application.

To remove applications from Region 1, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cfs region 4`	Enters the Region 4.
Step 3	`switch(config-cfs-region)# no ntp` `switch(config-cfs-region)# no callhome`	Removes application(s) that belong to Region 1, which you want to move.

Deleting CFS Regions

Deleting a region is nullifying the region definition. All the applications bound by the region are released back to the default region by deleting that region.

To delete a region, for example, a region numbered 4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# no cfs region 4`
	`WARNING: All applications in the region wiil be moved to default region.`
	`Are you sure? (y/n) [n]`
		Deletes the Region 4.

Note After Step 2, you see the warning, "All the applications in the region will be moved to the default region."

Default Settings

Table 6-1 lists the default settings for CFS configurations.

Table 6-1 Default CFS Parameters

Parameters	Default
CFS distribution on the switch	Enabled.
Database changes	Implicitly enabled with the first configuration change.
Application distribution	Differs based on application.
Commit	Explicit configuration is required.
CFS over IP	Disabled.
IPv4 multicast address	239.255.70.83
IPv6 multicast address	ff15::efff:4653

CHAPTER 7

Software Images

This chapter describes how to install and upgrade Cisco MDS SAN-OS software images. It includes the following sections:

About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The Cisco MDS SAN-OS consists of two images—the kickstart image and the system image. To upgrade the switch to a new image, you must specify the variables that direct the switch to the images.

- To select the kickstart image, use the KICKSTART variable.
- To select the system image, use the SYSTEM variable.

The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch. Both images are not always required for each install.

Note Unless explicitly stated, the software install procedures in this chapter apply to any switch in the Cisco MDS 9000 Family.

Dependent Factors for Software Installation

The software image install procedure is dependent on the following factors:

- Software images—The kickstart and system image files reside in directories or folders that can be accessed from the Cisco MDS 9000 Family switch prompt.
- Image version—Each image file has a version.
- Flash disks on the switch—The bootflash: resides on the supervisor module and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules.

Selecting the Correct Software Images for Cisco MDS 9100 Series Switches

The Supervisor-1 and Supervisor-2 modules supported by Cisco MDS 9100 Series switches require different system and kicstart images. You can determine which images to use on your switch by the naming conventions shown in Table 7-1.

Table 7-1 Supervisor Module Software Image Naming Conventions for MDS 9100 Series

Cisco MDS 9100 Series Switch Type	Supervisor Module Type	Naming Convention
9120 or 9140	Supervisor-1 module	Filename begins with m9100-s1ek9
9124, 9134, Cisco Fabric Switch for HP c-Class BladeSystem, Cisco Fabric Switch for IBM BladeCenter	Supervisor-2 module	Filename begins with m9100-s2ek9

Selecting the Correct Software Images for Cisco MDS 9200 Series Switches

The Supervisor-1 and Supervisor-2 modules supported by Cisco MDS 9200 Series switches require different system and kicstart images. You can determine which images to use on your switch by the naming conventions shown in Table 7-2.

Table 7-2 Supervisor Module Software Image Naming Conventions for MDS 9200 Series

Cisco MDS 9200 Series Switch Type	Supervisor Module Type	Naming Convention
9222i	Supervisor-2 module	Filename begins with m9200-s2ek9
9216, 9216A or 9216i	Supervisor-1 module	Filename begins with m9200-s1ek9

Selecting the Correct Software Images for Cisco MDS 9500 Family Switches

The Supervisor-1 and Supervisor-2 modules supported by Cisco MDS 9500 Family switches require different system and kickstart images. You can determine which images to use on your switch by the naming conventions shown in Table 7-3.

Table 7-3 Supervisor Module Software Image Naming Conventions for MDS 9500 Series

Cisco MDS 9500 Series Switch Type	Supervisor Module Type	Naming Convention
9506 or 9509	Supervisor-1 module	Filename begins with m9500-sf1ek9
	Supervisor-2module	Filename begins with m9500-sf2ek9
9513	Supervisor-2 module	Filename begins with m9500-sf2ek9

Use the **show module** command to display the type of supervisor module in the switch.

Example 7-1 shows the output for a switch with Supervisor-1 modules.

Example 7-1 show module Command Output for Supervisor-1 Modules

```
switch# show module
Mod Ports Module-Type Model Status
--- ----- -------------------------------- ------------------ ------------
...
...
5 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 active*
6 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 ha-standby
```

Example 7-3 shows the output for a switch with Supervisor-2 modules.

Example 7-2 show module Command Output for Supervisor-2 Modules

```
switch# show module
Mod Ports Module-Type Model Status
--- ----- -------------------------------- ------------------ ------------
...
...
7 0 Supervisor/Fabric-2 DS-X9530-SF2-K9 active *
8 0 Supervisor/Fabric-2 DS-X9530-SF2-K9 ha-standby
```

Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- Customer Service

  Before performing any software upgrade, contact your respective customer service representative to review your software upgrade requirements and to provide recommendations based on your current operating environment.

Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

- Scheduling

  Schedule the upgrade when the fabric is stable and steady. Ensure that everyone who has access to the switch or the network is not configuring the switch or the network during this time. All configurations are disallowed at this time.

- Space

  Verify that sufficient space is available in the location where you are copying the images. This location includes the active and standby supervisor module bootflash: (internal to the switch).

  - Standby supervisor module bootflash: file system (see Chapter 5, “Initial Configuration”).
  - Internal bootflash: offers approximately 200 MB of user space.

- Hardware

  Avoid power interruption during any install procedure. These kinds of problems can corrupt the software image.

- Connectivity (to retrieve images from remote servers)

  - Configure the IPv4 address or IPv6 address for the 10/100/1000 BASE-T Ethernet port connection (interface mgmt0).

Note 1000 BASE-T Ethernet is only available on Supervisor-2 modules.

  - Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router to route traffic between subnets.

- Images

  - Ensure that the specified system and kickstart images are compatible with each other.
  - If the kickstart image is not specified, the switch uses the current running kickstart image.
  - If you specify a different system image, ensure that it is compatible with the running kickstart image.
  - Retrieve images in one of two ways:

    Local file—images are locally available on the switch.

    Network file—images are in a remote location and the user specifies the destination using the remote server parameters and the file name to be used locally.

- Terminology

  Table 7-4 summarizes terms used in this chapter with specific reference to the install and upgrade process.

Table 7-4 Terms Specific to This Chapter

Term		Definition
bootable		The modules ability to boot or not boot based on image compatibility.
impact		The type of software upgrade mechanism—disruptive or nondisruptive.
install-type	reset	Resets the module.
	sw-reset	Resets the module immediately after switchover.
	rolling	Upgrades each module in sequence.
	copy-only	Updates the software for BIOS, loader, or bootrom.

- Commands
  - Verify connectivity to the remote serverusing the **ping** command.
  - Ensure that the required space is available for the image files to be copied using the **dir** command.
  - We recommend the one-step **install all** command to upgrade your software. This command upgrades all modules in any Cisco MDS 9000 Family switch (see the "Benefits of Using the install all Command" section on page 7-7).
  - Run only one installation on a switch at any time.
  - Do not issue another command while running the installation.
  - Do the installation on the active supervisor module, not the standby supervisor module.

Note If the switching module(s) are not compatible with the new supervisor module image, some traffic disruption may be noticed in the related modules, depending on your configuration. These modules are identified in the summary when you issue the **install all** command. You can choose to proceed with the upgrade or end at this point.

Note When you issue the **install all** command, the switch displays a summary of changes that are made to your configuration.

Note Prior to Cisco SAN-OS Release 3.0, to preserve the FC IDs in your configuration, verify that the persistent FC ID feature is enabled before rebooting. This feature is enabled by default. In earlier releases, the default is disabled. See the "FC IDs" section on page 17-14.

Software Upgrade Methods

You can upgrade software without any disruptions using the Cisco MDS SAN-OS software designed for mission-critical high availability environments. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9500 Directors, we highly recommend that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9000 Family using one of the following methods:

- Automated, one-step upgrade using the **install all** command. This upgrade is nondisruptive for directors in the Cisco MDS 9500 Series (see the "Automated Upgrades" section on page 7-7).
- Quick, one-step upgrade using the **reload** command. This upgrade is disruptive (see the "Quick Upgrade" section on page 7-31).

Tip The **install all** command compares and presents the results of the compatibility before proceeding with the installation. You can exit if you do not want to proceed with these changes.

In some cases, regardless of which process you use, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor module system with kickstart or system image changes.
- A dual supervisor module system with incompatible system software images.

Note For high availability, you need to connect the ethernet port for both active and standby supervisors to the same network or virtual LAN. The active supervisor owns the one IP address used by these ethernet connections. On a switchover, the newly activated supervisor takes over this IP address.

Determining Software Compatibility

If the running image and the image you want to install are incompatible, the software reports the incompatibility. In some cases, you may decide to proceed with this installation. If the active and the standby supervisor modules run different versions of the image, both images may be HA compatible in some cases and incompatible in others.

Compatibility is established based on the image and configuration:

- Image incompatibility—The running image and the image to be installed are not compatible.
- Configuration incompatibility—There is a possible incompatibility if certain features in the running image are turned off as they are not supported in the image to be installed. The image to be installed is considered incompatible with the running image if one of the following statements is true:
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and may cause the switch to move into an inconsistent state. In this case, the incompatibility is *strict.*
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and does not cause the switch to move into an inconsistent state. In this case, the incompatibility is *loose.*

To view the results of a dynamic compatibility check, issue the **show incompatibility system bootflash:***filename* command (see Example 7-3). Use this command to obtain further information when the **install all** command returns the following message:

```
Warning: The startup config contains commands not supported by the standby supervisor; as
a result, some resources might become unavailable after a switchover.

Do you wish to continue? (y/ n) [y]: n
```

Example 7-3 Displays HA Compatibility Status

```
switch# show incompatibility system bootflash:new-system-image
The following configurations on active are incompatible with the system image
The following configurations on active are incompatible with the system image
1) Service : cfs , Capability : CAP_FEATURE_CFS_ENABLED_DEVICE_ALIAS
Description : CFS - Distribution is enabled for DEVICE-ALIAS
Capability requirement : STRICT
Disable command : no device-alias distribute
```

Automated Upgrades

The **install all** command upgrades all modules in any Cisco MDS 9000 Family switch. Figure 7-1 provides an overview of the switch status before and after issuing the **install all** command.

Figure 7-1 The Effect of the install all Command

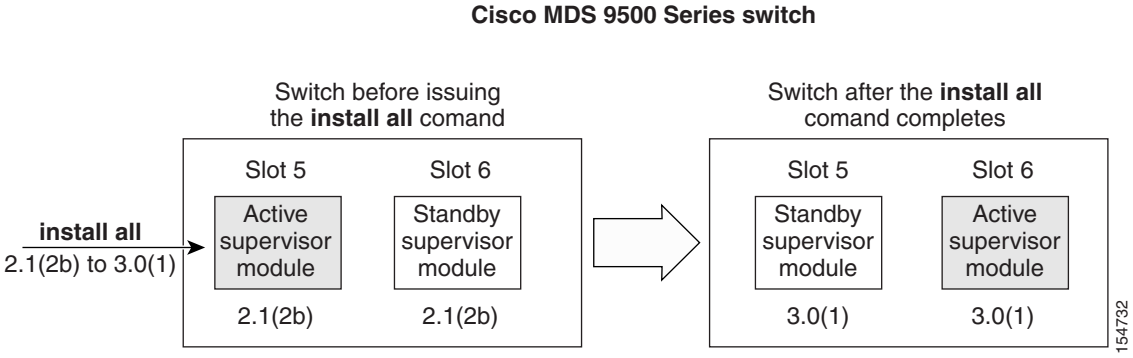

The **install all** command automatically verifies if the standby supervisor module is functioning (if present). If it is not functioning, it reloads that module and uses the **reload module** *slot* **force-dnld** command to force it to function.

Benefits of Using the install all Command

The **install all** command provides the following benefits:

- You can upgrade the entire switch using just one procedurecommand.
- You can receive descriptive information on the intended changes to your system before you continue with the installation.
- You have the option to cancel the command. Once the effects of the command are presented, you can continue or cancel when you see this question (the default is **no**):

  ```
  Do you want to continue (y/n) [n] :y
  ```

- You can upgrade the entire switch using the least disruptive procedure.

- You can see the progress of this command on the console, Telnet, and SSH screens:
  - After a switchover process, you can see the progress from both the supervisor modules.
  - Before a switchover process, you can only see the progress from the active supervisor module.
- The command automatically checks the image integrity. This includes the running kickstart and system images.
- The command performs a platform validity check to verify that a wrong image is not used—for example, to check if an MDS 9500 Series image is used inadvertently to upgrade an MDS 9200 Series switch.
- The **Ctrl-c** escape sequence gracefully ends the command. The command sequence completes the update step in progress and returns to the switch prompt. (Other upgrade steps cannot be ended using **Ctrl-c**.)
- After issuing the command, if any step in the sequence fails, the command completes the step in progress and ends.

  For example, if a switching module fails to be updated for any reason (for example, due to an unstable fabric state), then the command sequence disruptively updates that module and ends. In such cases, you can verify the problem on the affected switching module and upgrade the other switching modules.

Recognizing Failure Cases

The following situations cause the installation to end:

- If the standby supervisor module bootflash: file system does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the fabric or switch is configured while the upgrade is in progress.
- If the **install all** command is issued on the standby supervisor module.
- If the **install all** command does not reference the default bootflash: in a dual supervisor module configuration.
- If a module is removed while the upgrade is in progress.
- If the switch has any power disruption while the upgrade is in progress.
- If the entire path for the remote location is not specified accurately.
- If images are incompatible after an upgrade. For example, a switching module image may be incompatible with the system image, or a kickstart image may be incompatible with a system image. This is also identified by the **show install all impact** command in the compatibility check section of the output (under the Bootable column).

Caution If the **install all** command is ended, be sure to verify the state of the switch at every stage and reissue the command after 10 seconds. If you reissue the **install all** command within the 10-second span, the command is rejected with an error message indicating that an installation is currently in progress.

Tip All configurations are disallowed while the **install all** command is in progress. However, configurations coming through the CFS applications are allowed and may affect the upgrade procedure.

Using the install all Command

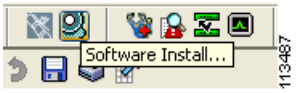

Note Ensure that there is enough space available on the active and standby supervisor module bootflash: to store the images being installed, even if the images are supplied in slot0:. The system will automatically synchronize the images to the standby supervisor module.

To perform an automated software upgrade on any switch, follow these steps:

Step 1 Log into the switch through the console, Telnet, or SSH port of the active supervisor module.

Step 2 Create a backup of your existing configuration file, if required (see the "Working with Configuration Files" section on page 8-1).

Step 3 Verify that you have enough free space available on the active and standby supervisor module bootflash:. The download site on Cisco.com shows the size of the system image file in bytes. If there is not adequate space, delete files using the **delete** *filename* EXEC command.

```
switch# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch# show module
Mod  Ports  Module-Type                        Model              Status
---  -----  ---------------------------------- ------------------ ------------
2    32     Storage Services Module            DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1                DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1                DS-X9530-SF1-K9    ha-standby
...
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch(standby)# exit
```

```
switch#
```

Step 4 Download a Cisco SAN-OS system image to the active supervisor module bootflash: from a TFTP server, if necessary.

```
switch# copy tftp://10.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server......
```

Note Ensure that you download the correct Cisco SAN-OS system image. The system image for Supervisor-1 modules is different from the system image for Supervisor-2 modules.

Step 5 Download a kickstart image to the active supervisor module bootflash: from a TFTP server, if necessary.

```
switch# copy tftp://10.16.10.100/kickstart-img bootflash:kickstart-img
Trying to connect to tftp server......
```

Step 6 Perform the upgrade by issuing the **install all** command.

Note On a dual supervisor module switch, always use the default bootflash: in the **install all** command syntax. Do not qualify it with "//sup-active/" or "//sup-local". Always use the following syntax: **install all system bootflash:***filename* **kickstart bootflash:***filename*.

```
switch# install all system bootflash:system-img kickstart bootflash:kickstart-img

Verifying image bootflash:/kickstart-img
[####################] 100% -- SUCCESS

Verifying image bootflash:/system-img
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/kickstart-img.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/kickstart-img.
[####################] 100% -- SUCCESS

Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes      disruptive       rolling  Hitless upgrade is not supported
     3       yes      disruptive       rolling  Hitless upgrade is not supported
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset

Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     1         slc               1.3(2a)                1.3(1)           yes
```

```
        1        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        2         ips              1.3(2a)               1.3(1)         yes
        2        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        3         ips              1.3(2a)               1.3(1)         yes
        3        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        4         slc              1.3(2a)               1.3(1)         yes
        4        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        5      system              1.3(2a)               1.3(1)         yes
        5  kickstart               1.3(2a)               1.3(1)         yes
        5        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        5      loader                1.2(2)               1.2(2)          no
        6      system              1.3(2a)               1.3(1)         yes
        6  kickstart               1.3(2a)               1.3(1)         yes
        6        bios     v1.1.0(10/24/03)     v1.1.0(10/24/03)          no
        6      loader                1.2(2)               1.2(2)          no

Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Syncing image bootflash:/kickstart-img to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:/system-img to standby.
[####################] 100% -- SUCCESS
Jan 18 23:40:03 Hacienda %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 6: Waiting for module online.
|
Auto booting bootflash:/kickstart-img bootflash:/system-img...
Booting kickstart image: bootflash:/kickstart-img....
....................................Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..r.r.. done.
Loading system software
Uncompressing system image: bootflash:/system-img
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3
```

Step 7 Exit the switch console and open a new terminal session to view the upgraded supervisor module using the **show module** command.

If the configuration meets all guidelines when the **install all** command is issued, all modules (supervisor and switching) are upgraded. This is true for any switch in the Cisco MDS 9000 Family.

Caution If a nondisruptive upgrade operation fails for any reason other than those listed in the "Recognizing Failure Cases" section on page 7-8, contact your reseller or Cisco representative for further assistance.

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Upgrading Services Modules

Any Fibre Channel switching module supports nondisruptive upgrades. The14/2-port Multiprotocol Services (MPS-14/2)) module supports nondisruptive upgrades for the Fibre Channel ports. Any software upgrade for the two Gigabit Ethernet ports in this module is disruptive. See Chapter 44, "Configuring IP Storage," for more information on MPS-14/2 modules.

Caution Any software upgrade for the Caching Services Module (CSM) and the IP Storage (IPS) services modules is disruptive.

CSMs and IPS modules use a rolling upgrade install mechanism to guarantee a stable state for each module in the switch:

- Each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded. See the Chapter 44, "Configuring IP Storage," for more information on IPS modules.
- Each CSM module requires a 30-minute delay before the next CSM module is upgraded. Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

When you upgrade, or downgrade, the SSI boot image on a Storage Services Module (SSM), you might disrupt traffic through the module. Table 7-5 describes how updating the SSI boot image affects SSM traffic.

Table 7-5 SSI Boot Image Updating Affects on SSM Traffic

Cisco MDS SAN-OS Release	Traffic Type	Disrupts Traffic?
2.0(2b) through 2.1(1a)	All	Yes
2.1(2) and later	Layer 2 Fibre Channel switching only	No
	Both Layer 2 Fibre Channel switching and Layer 3 Intelligent Storage Services (such as FCWA, NASB, SANTap, ISAPI virtualization)	Yes
	Layer 3 Intelligent Storage Services (such as FCWA, NASB, SANTap, ISAPI virtualization) only	Yes

As shown in Table 7-5, Layer 3 Intelligent Storage Services traffic is disrupted when you update the SSI boot image. If you have configured Layer 3 Intelligent Storage Services on your SSM, we recommend that you shut down these services before upgrading the SSI boot image. You can use dual fabric configuration to minimize the impact of shutting down Layer 3 services. See Chapter 11, "Managing Modules," for more information on updating the boot images on the SSM.

Sample install all Commands

Example 7-4 displays the result of the **install all** command issued from a console terminal that is connected to the active supervisor module. Once a switchover happens, you can see the rest of the output from the console terminal of the standby supervisor module. Example 7-5 displays the file output continuation of the **install all** command on the console of the standby supervisor module. Example 7-6 displays the result of the **install all** command issued from a console terminal for a system that contains an SSI image.

Similarly, you can view the results of the **install all** command issued from the SSH or Telnet terminal that is connected to the active supervisor module. Once a switchover happens, you need to log back into the switch and issue the **show install all status** command (see the "Upgrade Status Verification" section on page 7-20).

Example 7-4 Successful install all Command Issued from the Active Console

```
Hacienda# install all system bootflash:system-img kickstart bootflash:kickstart-img

Verifying image bootflash:/kickstart-img
[####################] 100% -- SUCCESS

Verifying image bootflash:/system-img
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/system-img.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/kickstart-img.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/kickstart-img.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes      disruptive       rolling  Hitless upgrade is not supported
     3       yes      disruptive       rolling  Hitless upgrade is not supported
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset



Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     1         slc               1.3(2a)               1.3(1)           yes
     1        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     2         ips               1.3(2a)               1.3(1)           yes
     2        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     3         ips               1.3(2a)               1.3(1)           yes
```

```
     3       bios     v1.1.0(10/24/03)      v1.1.0(10/24/03)             no
     4        slc                1.3(2a)                1.3(1)          yes
     4       bios     v1.1.0(10/24/03)      v1.1.0(10/24/03)             no
     5     system                1.3(2a)                1.3(1)          yes
     5 kickstart                 1.3(2a)                1.3(1)          yes
     5       bios     v1.1.0(10/24/03)      v1.1.0(10/24/03)             no
     5     loader                 1.2(2)                1.2(2)             no
     6     system                1.3(2a)                1.3(1)          yes
     6 kickstart                 1.3(2a)                1.3(1)          yes
     6       bios     v1.1.0(10/24/03)      v1.1.0(10/24/03)             no
     6     loader                 1.2(2)                1.2(2)             no

Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Syncing image bootflash:/kickstart-img to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:/system-img to standby.
[####################] 100% -- SUCCESS
Jan 18 23:40:03 Hacienda %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 6: Waiting for module online.
|
Auto booting bootflash:/kickstart-img bootflash:/system-img...
Booting kickstart image: bootflash:/kickstart-img....
.....................................Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..r.r.. done.
Loading system software
Uncompressing system image: bootflash:/system-img
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3
```

Example 7-5 displays the file output continuation of the **install all** command on the console of the standby supervisor module.

Example 7-5 Successful install all Command Output Continued from the Standby Console

```
Hacienda(standby)#

Auto booting bootflash:/kickstart-img bootflash:/system-img...
Booting kickstart image: bootflash:/kickstart-img....
.....................................Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..r.r.. done.
Loading system software
Uncompressing system image: bootflash:/system-img
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
CCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3
```

```
Continue on installation process, please wait.
The login will be disabled until the installation is completed.

Module 6: Waiting for module online.
Jan 18 23:43:02 Hacienda %PORT-5-IF_UP: Interface mgmt0 is up
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
FM_SERVER_PKG. Application(s) shutdown in 53 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
ENTERPRISE_PKG. Application(s) shutdown in 50 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
SAN_EXTN_OVER_IP. Application(s) shutdown in 50 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LICAPP_NO_LIC: Application port-security running
without ENTERPRISE_PKG license, shutdown in 50 days
Jan 18 23:43:19 Hacienda %LICMGR-4-LOG_LICAPP_EXPIRY_WARNING: Application Roles evaluation
license ENTERPRISE_PKG expiry in 50 days
Jan 18 23:44:54 Hacienda %BOOTVAR-5-NEIGHBOR_UPDATE_AUTOCOPY: auto-copy supported by
neighbor, starting...
Module 1: Non-disruptive upgrading.
[#                   ]   0%Jan 18 23:44:56 Hacienda %MODULE-5-STANDBY_SUP_OK: Supervisor 5
is standby
Jan 18 23:44:55 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_STARTED:  Module image download
process. Please wait until completion...
Jan 18 23:45:12 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_COMPLETE:  Module image download
process. Download successful.
Jan 18 23:45:48 Hacienda %MODULE-5-MOD_OK: Module 1 is online
[####################] 100% -- SUCCESS
Module 4: Non-disruptive upgrading.
[#                   ]   0%Jan 18 23:46:12 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_STARTED:
Module image download process. Please wait until completion...
Jan 18 23:46:26 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_COMPLETE:  Module image download
process. Download successful.
Jan 18 23:47:02 Hacienda %MODULE-5-MOD_OK: Module 4 is online
[####################] 100% -- SUCCESS
Module 2: Disruptive upgrading.
...
-- SUCCESS
Module 3: Disruptive upgrading.
...
 -- SUCCESS
Install has been successful.
MDS Switch
Hacienda login:
```

Example 7-6 displays the result of the **install all** command issued from a console terminal for a system that contains an SSI image.The **install all** command syncs the SSI image to the standby supervisor module.

Note You can use the **install all** command for the SSM only if the SSM is already up and running. For first time SSM installations, see the "Upgrading the SSI Boot Image on an SSM" section on page 11-19.

Example 7-6 Successful install all Command Including an SSI Image

```
Cisco-MDS# install all system bootflash:isan-2-1-1a kickstart
bootflash:boot-2-1-1a  ssi bootflash:ssi-2.1.1a

Verifying image bootflash:/ssi-2.1.1a
[####################] 100% -- SUCCESS

Verifying image bootflash:/boot-2-1-1a
```

```
[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-2-1-1a
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-2-1-1a.
[####################] 100% -- SUCCESS

Extracting "ips4" version from image bootflash:/isan-2-1-1a.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-2-1-1a.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-2-1-1a.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/boot-2-1-1a.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     2       yes  non-disruptive       rolling
     3       yes      disruptive       rolling  Hitless upgrade is not supported
     4       yes      disruptive       rolling  Hitless upgrade is not supported
     5       yes  non-disruptive         reset



Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     2         slc                2.0(3)                2.1(1a)           yes
     2        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     3         slc                2.0(3)                2.1(1a)           yes
     3         ssi                2.0(3)                2.1(1a)           yes
     3        bios      v1.0.8(08/07/03)      v1.1.0(10/24/03)           yes
     4        ips4                2.0(3)                2.1(1a)           yes
     4        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     5      system                2.0(3)                2.1(1a)           yes
     5   kickstart                2.0(3)                2.1(1a)           yes
     5        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     5      loader                1.2(2)                 1.2(2)            no

Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Module 6:Force downloading.
 -- SUCCESS

Syncing image bootflash:/ssi-2.1.1a to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:/boot-2-1-1a to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:/isan-2-1-1a to standby.
[####################] 100% -- SUCCESS

Setting boot variables.
```

```
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 3:Upgrading Bios/loader/bootrom.
[####################] 100% -- SUCCESS

Module 6:Waiting for module online.
-- SUCCESS


"Switching over onto standby".
----------------------------
```

Note If you perform the **install all** command to downgrade to a Cisco MDS SAN-OS release that does not support the SSM module, you must power down the SSM module when prompted by the CLI console. The boot variables for the SSM module are lost.

Example 7-7 displays the result of the **install all** command if the system and kickstart files are automatically downloaded using a remote (TFTP, FTP, SCP, or SFTP) download option. It shows an accurate and complete example.

Caution Specify the complete path of the remote location. The system will not allow you to proceed if the entire path is not accurately specified. Here are examples of incomplete **install all** commands.

```
switch# install all system bootflash:system-image kickstart tftp:
Please provide a complete URI
switch# install all system scp:
Please provide a complete URI
```

Example 7-7 A Sample of the install all Command Issued Using a Remote Download

```
switch# install all system
scp://user@10.10.1.1/tftpboot/HKrel/qa/final/m9500-sf1ek9-mz.1.3.2a.bin kickstart
scp://user@10.10.1.1/tftpboot/HKrel/qa/final/m9500-sf1ek9-kickstart-mz.1.3.2a.bin
For scp://user@10.10.1.1, please enter password:
For scp://user@10.10.1.1, please enter password:

Copying image from
scp://user@10.10.1.1/tftpboot/HKrel/qa/final/m9500-sf1ek9-kickstart-mz.1.3.2a.bin to
bootflash:///m9500-sf1ek9-kickstart-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS

Copying image from scp://user@10.10.1.1/tftpboot/HKrel/qa/final/m9500-sf1ek9-mz.1.3.2a.bin
to bootflash:///m9500-sf1ek9-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-kickstart-mz.1.3.2a.bin
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-mz.1.3.2a.bin
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sf1ek9-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sf1ek9-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:///m9500-sf1ek9-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:///m9500-sf1ek9-kickstart-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sf1ek9-kickstart-mz.1.3.2a.bin.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes      disruptive       rolling  Hitless upgrade is not supported
     3       yes  non-disruptive       rolling
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset
     7       yes  non-disruptive       rolling
     8       yes  non-disruptive       rolling
     9       yes      disruptive       rolling  Hitless upgrade is not supported


Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     1         slc                1.3(1)                1.3(2a)           yes
     1        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     2         ips                1.3(1)                1.3(2a)           yes
     2        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     3         slc                1.3(1)                1.3(2a)           yes
     3        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     4         slc                1.3(1)                1.3(2a)           yes
     4        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     5      system                1.3(1)                1.3(2a)           yes
     5   kickstart                1.3(1)                1.3(2a)           yes
     5        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     5      loader                1.2(2)                 1.2(2)            no
     6      system                1.3(1)                1.3(2a)           yes
     6   kickstart                1.3(1)                1.3(2a)           yes
     6        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     6      loader                1.2(2)                 1.2(2)            no
     7         slc                1.3(1)                1.3(2a)           yes
     7        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     8         slc                1.3(1)                1.3(2a)           yes
     8        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     9         ips                1.3(1)                1.3(2a)           yes
     9        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no

Do you want to continue with the installation (y/n)?  [n]
```

Example 7-8 displays the **install all** command output of a failed operation due to a lack of disk space.

Example 7-8 Failed Operation Due to a Full bootflash: File System

```
switch# install all system bootflash:isan-1.3.2a kickstart bootflash:boot-1.3.2a

Verifying image bootflash:/boot-1.3.2a
```

```
[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-1.3.2a
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-1.3.2a.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:/isan-1.3.2a.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-1.3.2a.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-1.3.2a.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/boot-1.3.2a.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes      disruptive       rolling  Hitless upgrade is not supported
     3       yes  non-disruptive       rolling
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset
     7       yes  non-disruptive       rolling
     8       yes  non-disruptive       rolling
     9       yes      disruptive       rolling  Hitless upgrade is not supported


Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     1         slc                1.3(1)                1.3(2a)           yes
     1        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     2         ips                1.3(1)                1.3(2a)           yes
     2        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     3         slc                1.3(1)                1.3(2a)           yes
     3        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     4         slc                1.3(1)                1.3(2a)           yes
     4        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     5      system                1.3(1)                1.3(2a)           yes
     5   kickstart                1.3(1)                1.3(2a)           yes
     5        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     5      loader                1.2(2)                 1.2(2)            no
     6      system                1.3(1)                1.3(2a)           yes
     6   kickstart                1.3(1)                1.3(2a)           yes
     6        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     6      loader                1.2(2)                 1.2(2)            no
     7         slc                1.3(1)                1.3(2a)           yes
     7        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     8         slc                1.3(1)                1.3(2a)           yes
     8        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no
     9         ips                1.3(1)                1.3(2a)           yes
     9        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)            no

Do you want to continue with the installation (y/n)?  [n] y
```

```
Install is in progress, please wait.

Syncing image bootflash:/boot-1.3.2a to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:/isan-1.3.2a to standby.
[#                   ]   0% -- FAIL. Return code 0x401E0008 (request was aborted, standby
disk may be full).

Install has failed. Return code 0x40930013 (Syncing images to standby failed).
Please identify the cause of the failure, and try 'install all' again.
Dec 15 19:36:42 switch %SYSMGR-3-SERVICE_TERMINATED: Service "installer" (PID 5470) has
finished with error code SYSMGR_EXITCODE_FAILURE_NOCALLHOME (20).
```

Example 7-9 displays the **install all** command output of a failed operation due to an invalid image.

Example 7-9 Failed Operation Due to an Invalid Image

```
install all system bootflash:junk kickstart bootflash:junk

Verifying image bootflash:/junk
[#                   ]   0% -- FAIL. Return code 0x4045001E (mismatch between actual image
type and boot variable).
Compatibility check failed. Return code 0x40930011 (Image verification failed).
Hacienda# Jan 19 00:20:35 Hacienda %SYSMGR-3-SERVICE_TERMINATED: Service "installer" (PID
5664) has finished with error code SYSMGR_EXITCODE_FAILURE_NOCALLHOME (20).
```

Upgrade Status Verification

Use the **show install all status** command to view the ongoing **install all** command or the log of the last installed **install all** command from a console, SSH, or Telnet session.

This command presents the **install all** output on both the active and standby supervisor module even if you are not connected to the console terminal. It only displays the status of an **install all** command that is issued from the CLI (not the GUI). See Example 7-10.

Example 7-10 Displays the install all Command Output

```
switch# show install all status
There is an on-going installation... <---------------------- in progress installation
Enter Ctrl-C to go back to the prompt.
Verifying image bootflash:/b-1.3.0.104
 -- SUCCESS

Verifying image bootflash:/i-1.3.0.104
 -- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
 -- SUCCESS

Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
 -- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
 -- SUCCESS
```

```
switch# show install all status
This is the log of last installation.          <<<<<< log of last install

Verifying image bootflash:/b-1.3.0.104
 -- SUCCESS

Verifying image bootflash:/i-1.3.0.104
 -- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
 -- SUCCESS

Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
 -- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
 -- SUCCESS
```

Non-Disruptive Upgrades on Fabric and Modular Switches

This section describes how to perform non-disruptive upgrades on the following Cisco Fabric Switches:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Multilayer Fabric Switch
- Cisco MDS 9222i Multiservice Modular Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

This section includes the following topics:

Preparing for a Non-Disruptive Upgrade on Fabric and Modular Switches

You can upgrade software on the following without any disruptions using the **install all** command for the system software images.

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Multlayer Fabric Switch
- Cisco MDS 9222i Multiservice Modular Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

When completed, the supervisor kickstart image, supervisor system image, the linecard image and the system bios are all updated.

Non-disruptive upgrades on these fabric switches take down the control plane for not more than 80 seconds. In some cases, when the upgrade has progressed past the point at which it cannot be stopped gracefully, or if a failure occurs, the software upgrade may be disruptive.

Note During the upgrade the control plane is down, but the data plane remains up. So new devices will be unable to log in to the fabric via the control plane, but existing devices will not experience any disruption of traffic via the data plane.

Before attempting to upgrade any software images on these fabric switches, follow these guidelines:

- During the upgrade, the fabric must be stable. None of the following configuration activities are allowed:
  - zoning changes
  - telnet sessions
  - schedule changes
  - switch cabling
  - addition or removal of physical devices
- Configure the FSPF timers to the default value of 20 seconds.
- If there are any CFS commits pending in the fabric, the upgrade is aborted.
- If there is a zone server merge in progress, the upgrade is aborted.
- If the upgrade is aborted due to a service not being ready for the upgrade, you are prompted to enter the **show install all failure-reason** command to identify the reason why.
- If there is insufficient space available in the system to load the new images, then you will be notified via the compatibility table. At this point, you need to either abort the upgrade or proceed with a disruptive upgrade.
- Issue the **no logging level all** command before the commencing with the upgrade. Failing to issue this command could result in a failure due to the debug system log messages being printed and potentially resulting in control plane downtime exceeding 80 seconds.
- If VRRP is running on the mgmt0 interface, and the switch being upgraded is the master, then a new master is selected. This cannot be avoided because the mgmt0 interface goes down when the control plane goes down.
- On the Cisco MDS 18/4-port multiservice module, upgrades of the 4 Gigabit Ethernet ports for the hybrid Supervisor 18/4 line card will be disruptive.

To ensure that you can view the entire upgrade process, it is recommended that you perform the upgrade via the console port; performing the upgrade in this way also enables you to log your session to a file (in case you need it later for troubleshooting). Also, telnet sessions are lost when the switch is rebooted, so if you wish to view the process in its entirety, use the console port instead.

Example 7-11 Failed Nondisruptive Upgrade Due to Insufficient Resources

```
switch# install all kickstart bootflash:boot-fs9124 system bootflash:isan-164

Verifying image bootflash:/boot-fs9124 for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-164 for boot variable "system".
[####################] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/isan-164.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-fs9124.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/isan-164.
[####################] 100% -- SUCCESS

Compatibility check is done:
Module  bootable   Impact         Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes   disruptive     reset  insufficient resources<----Reason for failure


Images will be upgraded according to following table:
Module  Image       Running-Version(pri:alt)          New-Version  Upg-Required
------  ----------  ------------------------          -----------  --------------------
     1      system  3.1(1u)                           3.1(1)           yes
     1   kickstart  3.1(1u)                           3.1(1)           yes
     1        bios  v1.0.0(10/04/06):v1.0.0(10/04/06) v1.0.0(10/04/06) no

Do you want to continue with the installation (y/n)?  [n]
```

Before performing an upgrade, you may wish to use the **show install all impact** command to view the effect of updating the system from the running image to another specified image.

```
switch# show install all impact kickstart bootflash:boot-fs9124 system bootflash:isan-164

Verifying image bootflash:/boot-fs9124 for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-164 for boot variable "system".
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-164.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-fs9124.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/isan-164.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive         reset



Images will be upgraded according to following table:
Module       Image   Running-Version(pri:alt)           New-Version      Upg-Required
------  ----------   ------------------------           -----------      -------------
     1      system   3.1(1u)                            3.1(1)           yes
     1   kickstart   3.1(1u)                            3.1(1)           yes
     1        bios   v1.0.0(10/04/06): v1.0.0(10/04/06) v1.0.0(10/04/06) no

switch#
```

Performing a Non-Disruptive Upgrade on a Fabric Switch

To perform a non-disruptive software upgrade on any of the following switches, enter the **install all kickstart** command using the console port:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Multlayer Fabric Switch
- Cisco MDS 9222i Multiservice Modular Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

```
switch# install all kickstart bootflash:boot-fs9124 system bootflash:isan-164u

Verifying image bootflash:/boot-fs9124 for boot variable "kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-164u for boot variable "system".
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-164u.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-fs9124.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/isan-164u.
[####################] 100% -- SUCCESS

Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive         reset

Images will be upgraded according to following table:
Module       Image  Running-Version(pri:alt)          New-Version           Upg-Required
------  ----------  -------------------------        --------------------  ------------
     1      system  3.1(1)                           3.1(1u)                yes
     1   kickstart  3.1(1)                           3.1(1u)                yes
     1        bios  v1.0.0(10/04/06): v1.0.0(10/04/06) v1.0.0(10/04/06)       no

Do you want to continue with the installation (y/n)?  [n]

Install is in progress, please wait.

Notifying services about the upgrade.
[####################] 100% -- SUCCESS

Setting boot variables.
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Converting startup config.
[####################] 100% -- SUCCESS
```

```
Upgrade can no longer be aborted, any failure will result in a disruptive upgrade.
<---Note that after this point you cannot abort the upgrade.

Freeing memory in the file system.
[####################] 100% -- SUCCESS

Loading images into memory.
[####################] 100% -- SUCCESS

Saving linecard runtime state.
[####################] 100% -- SUCCESS

Saving supervisor runtime state.
[####################] 100% -- SUCCESS

Saving mts state.
[####################] 100% -- SUCCESS

Rebooting the switch to proceed with the upgrade.


Continuing with installation process, please wait.
The login will be disabled until the installation is completed.

Status for linecard upgrade.
[####################] 100% -- SUCCESS

Performing supervisor state verification.
[####################] 100% -- SUCCESS

Install has been successful.
```

Caution It is recommended that you enable port-fast on the ethernet interface of the catalyst to which the management interface of the fabric switch is connected. This is to avoid spanning-tree convergence time on the catalyst and packets from the fabric switch are forwarded immediately during the non-disruptive upgrade.

Note When selecting images during the upgrade, ASM-SFN and SSI are not supported on the Cisco MDS 9124 Switch and the Cisco MDS 9134 Multilayer Fabric Switch.

Viewing the Status of a Non-Disruptive Upgrade on a Fabric Switch

You can view the status of a non-disruptive upgrade using the **show install all status** command. Note that the output shows the status only after the switch has rebooted to come up with the new image. All actions preceding the reboot are not captured in this output because when you enter the **install all** command using a telnet session, the session is disconnected when the switch reboots. After you can reconnect to the switch via a telnet session, the upgrade may already be complete; in this case, the output will show the status of the upgrade.

```
switch# show install all status
This is the log of last installation.

Continuing with installation process, please wait.
The login will be disabled until the installation is completed.
```

```
Status for linecard upgrade.
 -- SUCCESS

Performing supervisor state verification.
 -- SUCCESS

Install has been successful.
```

Troubleshooting a Non-Disruptive Upgrade on a Fabric Switch

When a non-disruptive upgrade begins, the system notifies all services that an upgrade is about to start, and finds out whether or not the upgrade can proceed. If a service cannot allow the upgrade to proceed at this time (for example, FSPF timers are not configured to the default value, or a CFS operation is in progress), then the service will abort the upgrade. In such a context, you will be prompted to enter the **show install all failure-reason** command to determine the reason why the upgrade cannot proceed.

```
...
Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Notifying services about the upgrade.
[#                   ]   0% -- FAIL. Return code 0x401E0066 (request timed out).

Please issue "show install all failure-reason" to find the cause of the failure.<---system
prompt to enter the show all failure-reason command.

Install has failed. Return code 0x401E0066 (request timed out).
Please identify the cause of the failure, and try 'install all' again.

switch# show install all failure-reason
Service: "cfs" failed to respond within the given time period.
switch#
```

If there are any failures for whatever reason (a save runtime state failure or linecard upgrade failure) once the upgrade is already in progress, then the switch will be rebooted disruptively because the changes cannot be rolled back. In such cases the upgrade has failed; you are not prompted to enter the **show install all failure-reason** command, and entering it will not yield any useful information.

If further assistance is required to determine why an upgrade is unsuccessful, you should collect the details from the **show tech-support** command output, and the console output from the installation, if available.

Manual Upgrade on a Dual Supervisor Module Switch

Caution If you are a new user, use the **install all** command to perform a software upgrade. This section is for administrators or individuals who are completely familiar with specific switch functions.

You can manually upgrade the BIOS and the loader in any Cisco MDS switch using the procedures provided in this section. This upgrade process requires you to implement some or all procedures depending on your switch or network configuration.

This section includes the following topics:

Preparing for a Manual Installation

To prepare any Cisco MDS 9000 Family switch for a manual software installation, follow these steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Create a backup of your existing configuration file, if required (see the "Saving the Running Configuration" section on page 8-4).

Step 3 Copy the software image from an SCP location to one of two targets: bootflash: or slot0:.

The switch remains operational while the image file is copied.

- Bootflash device (SCP defaults to the bootflash device)—Copy the software image file from the appropriate SCP file system to the bootflash: file system.

```
switch# copy scp://server_IP_address/destination_file_name
```

For example:

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```

Note The Cisco MDS 9216 Switch does not have an external CompactFlash (see the "Working with Configuration Files" section on page 8-1). If you are using a switch in this series, use the bootflash: file system to copy and verify files.

- CompactFlash device—Copy the software image file from the appropriate SCP file system to the CompactFlash device in slot0: file system.

```
switch# copy scp://server_IP_address/file_name_in_SCP slot0:system-image
```

You can also copy the image onto a new Flash disk from a PC and insert it in slot0: in the Cisco MDS 9500 Series switch. After you copy the image and insert it into the slot0: file system, the process is the same as the CompactFlash device after the **copy** command is issued.

Step 4 Verify that the file was copied in the required directory.

```
switch# dir bootflash:
40295206     Aug 05 15:23:51 1980  ilc1.bin
12456448     Jul 30 23:05:28 1980  kickstart-image1
12288        Jun 23 14:58:44 1980  lost+found/
27602159     Jul 30 23:05:16 1980  system-image1
12447232     Aug 05 15:08:30 1980  kickstart-image2
28364853     Aug 05 15:11:57 1980  system-image2

Usage for bootflash://sup-local
  135404544 bytes used
   49155072 bytes free
  184559616 bytes total
```

Step 5 Ensure that the software images are not damaged or corrupted in the saved bootflash: file system.

When copying a new image to your switch, confirm that the image was not corrupted during the copy process.

Use the **show version image** command to verify that the required image was copied successfully.

```
switch# show version image bootflash:kickstart-image
  image name: m9500-sf1ek9-kickstart-mzg.1.0.3.bin
  kickstart:  version 1.0(3)
  loader:     version 1.0(3)
  compiled:   2/12/2003 11:00:00
```

Note A verification failed message is generated when you use a Cisco MDS 9500 Series image on a Cisco MDS 9200 Series switch or a Cisco MDS 9200 Series image on a Cisco MDS 9500 Series switch. Be sure to verify the right image.

Step 6 Compare the running system image and the new image by issuing the **show install all impact** command.

Upgrading a Loader

The **install module** *slot# of the supervisor module* **loader** command upgrades the (boot) loader.

Note If the loader is upgraded, you need to reboot to make the new loader effective. You can schedule the reboot at a convenient time so traffic is not impacted.

Caution Before issuing this command, be sure to read the release notes to verify compatibility issues between the loader and the kickstart or system images.

To upgrade the loader on either the active or standby supervisor module, follow these steps.

Step 1 Use the **show version** command to verify the version on the active and standby supervisor modules.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 1.0.8
  loader:    version 1.1(2) <-------------------------current running version
  kickstart: version 2.0(1)
  system:    version 2.0(1)

  BIOS compile time:       08/07/03
  kickstart image file is: bootflash:///m9500-sf1ek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time:  10/25/2010 12:00:00
```

```
  system image file is:    bootflash:///m9500-sf1ek9-mzg.2.0.0.6.bin
  system compile time:     10/25/2020 12:00:00


Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:           0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
    Reason: Reset Requested by CLI command reload
    System version: 2.0(0.6)
    Service:
```

Step 2 Issue the **install module** command for the required supervisor module (active or standby). This example displays the command being issued for the standby supervisor module in slot 6.

```
switch# install module 6 loader bootflash:kickstart-image
```

Note If you install a loader version that is the same as the currently installed version, the command will not execute. When both the current version and the installed version are the same, use the **init system** command to force a loader upgrade.

Step 3 Use the **show version** command to verify the updated image on the supervisor module.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 1.3.1
  loader:    version 1.2(2) <---------------New running version
  kickstart: version 1.3(1) ]
  system:    version 1.3(1)

  BIOS compile time:       08/07/03
  kickstart image file is: bootflash:///m9500-sf1ek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time:  10/25/2010 12:00:00
  system image file is:    bootflash:///m9500-sf1ek9-mzg.2.0.0.6.bin
  system compile time:     10/25/2020 12:00:00


Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:           0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
```

```
    Reason: Reset Requested by CLI command reload
    System version: 2.0(0.6)
    Service:
```

Upgrading the BIOS

Tip Refer to the release notes to verify if the BIOS has changed for the image version being used.

Program the supervisor module or switching module BIOS only if a new BIOS image is provided by Cisco Systems. Only use the provided image to upgrade the BIOS. This command does not affect traffic and can be issued at any time on any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

Note If the BIOS is upgraded, reboot to make the new BIOS take effect. You can schedule the reboot at a convenient time so traffic is not impacted.

Caution The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To upgrade the BIOS for a module, follow these steps:

Step 1 Use the **show version** command to verify the current running BIOS version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license.
Software
  BIOS:      version 1.0(6) <---------- current running version
  loader:    version 1.0(3)
  kickstart: version 1.0(3)
  system:    version 1.0(3)

  BIOS compile time:       01/27/03
  kickstart image file is: bootflash:/kickstart-image
  kickstart compile time:  01/25/2003 12:00:00
  system image file is:    bootflash:/system-image
  system compile time:     01/25/2003 12:00:00

Hardware
  RAM 1027564 kB
```

Step 2 Verify that the BIOS version of the system image is different from the running image.

```
switch# show version image bootflash:system-image
  image name: m9500-sf1ek9-mz.1.0.3.bin
  bios:       version   v1.0.6(01/27/03) <---------- BIOS is same version 1.0.6
  system:     version 1.0(3)
  compiled:   2/28/2003 5:00:00
```

```
system service's list

package name             package version
acl                      1.0(3)
ascii-cfg                1.0(3)
bios_daemon              1.0(3)
...
```

Note If the versions are different, issue the **install module** command as specified in Step 3. If they are the same, you do not need to update the BIOS image.

Step 3 Run the **install module** *slot#* **bios** command to install each module (if required). In this example, the supervisor module in slot 6 was updated.

```
switch# install module 6 bios system bootflash:system-image
Started bios programming .... please wait
[########################################] 100%
BIOS upgrade succeeded for module 1
```

Caution Do not reboot the switch if any errors were indicated in response to this command.

Step 4 Issue the **show version** command to verify that the module was updated with a the new BIOS version.

```
switch# show version module 6
ModNo   Image Type   SW Version   SW Interim Version  BIOS Version
6       Stby Sup     1.3(2)       1.3(1.1)            1.1.0 [last 1.0.6]
```

Quick Upgrade

To perform a quick upgrade on a Cisco MDS 9000 Family switch, follow these steps:

Step 1 Copy the kickstart and system image files to the required location (see the "Copying Configuration Files" section on page 8-5).

Step 2 Set the boot variables.

```
switch# config t
switch(config)# boot system system-img
switch(config)# boot kickstart kickstart-img
switch(config)# exit
switch#
```

Step 3 Issue the **reload** command. The **reload** command reboots the system. This upgrade is disruptive.

Tip Use the **install all** command to gracefully reload the switch and handle configuration conversions.

Downgrading from a Higher Release

Use the **install all** command to gracefully reload the switch and handle configuration conversions. When downgrading any switch in the Cisco MDS 9000 Family, avoid using the **reload** command.

Note If you downgrade from Cisco MDS SAN-OS Release 3.1(3) to any earlier SAN-OS release after you execute the system default switchport mode F command, the ports retain the configuration that resultedfrom the execution of the command. In other words, the ports do not revert back to the mode they were in prior to executing the command.

For example, to revert to Cisco MDS SAN-OS Release 1.3(4b) or 1.3(5) from Release 3.0(1), follow these steps:

Step 1 Issue the **show incompatibility system** *image-filename* command to determine if you need to disable any features not supported by the older release. The command output provides the commands needed to disable the incompatible features.

```
switch# show incompatibility system bootflash:m9200-ek9-mz.1.3.4b.bin
The following configurations on active are incompatible with the system image
1) Service : cfs , Capability : CAP_FEATURE_CFS_ENABLED_DEVICE_ALIAS
Description : CFS - Distribution is enabled for DEVICE-ALIAS
Capability requirement : STRICT
Disable command : no device-alias distribute
```

Step 2 Disable the incompatible features.

```
switch# config t
switch(config)# no device-alias distribute
switch(config)# exit
switch#
```

Step 3 Save the running configuration to the startup configuration.

```
switch# copy running-config start-config
```

Step 4 Issue the **install all** command to downgrade the software (see the "Automated Upgrades" section on page 7-7).

Maintaining Supervisor Modules

This section includes general information about replacing and using supervisor modules effectively.

This section includes the following topics:

Replacing Supervisor Modules

To avoid packet loss when removing a supervisor module from a Cisco MDS 9500 Series Director, take the supervisor modules out of service before removing the supervisor module.

Use the out-of-service command in EXEC mode before removing the supervisor module.

out-of-service module *slot*

Where *slot* indicates the chassis slot number in which the supervisor module resides.

Note You must remove and reinsert or replace the supervisor module to bring it into service.

Migrating from Supervisor-1 Modules to Supervisor-2 Modules

Supervisor-1 modules and Supervisor-2 modules cannot be used in the same switch, except for migration purposes. Both the active and standby supervisor modules must be of the same type, either Supervisor-1 or Supervisor-2 modules. For Cisco MDS 9513 Directors, both supervisor modules must be Supervisor-2 modules.

The procedure described in this section ensures that your configuration is correctly synchronized after completing the migration.

Caution Migrating your supervisor modules is a disruptive operation.

Note Migrating from Supervisor-2 modules to Supervisor-1 modules is not supported.

To migrate from Supervisor-1 modules to Supervisor-2 modules on a Cisco MDS 9509 or 9506 switch, follow these steps:

Step 1 Ensure that the configured domain ID is the same as the current domain ID for every VSAN on the switch by following these steps:

a. Issue a **show vsan** command to display all the VSANs on the switch.

```
switch# show vsan
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 2 information
         name:VSAN0002  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 10 information
         name:VSAN0010  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id
         operational state:down

vsan 4094:isolated_vsan
```

b. Display the current and configured domain IDs for a VSAN.

```
switch# show fcdomain vsan 1
The local switch is the Principal Switch.

Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:05:30:00:35:df
        Running fabric name: 20:01:00:05:30:00:35:df
        Running priority: 128
        Current domain ID: 0x6a(106)

Local switch configuration information:
        State: Enabled
        FCID persistence: Enabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)

Principal switch run time information:
        Running priority: 128
```

c. Change the configured domain ID if it differs from the current domain ID.

```
switch# config t
switch(config)# fcdomain domain 106 static vsan 1
switch(config)# exit
switch#
```

d. Repeat Step b and Step c for each VSAN on the switch.

Step 2 Save the configuration.

```
switch# copy running-config startup-config
```

Step 3 Verify that the switch is running Cisco SAN-OS Release 3.0(1) or later. Upgrade the switch, if necessary (see the "Automated Upgrades" section on page 7-7).

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2005, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 0.0.11
  kickstart: version 3.0(1) [build 3.0(0.294)] [gdb]
  system:    version 3.0(1) [build 3.0(0.294)] [gdb]
...
```

Step 4 Issue a **show module** command to determine which Supervisor-1 module is the standby.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module                DS-X9016           ok
2    32     Storage Services Module           DS-X9032-SSM       ok
3    8      IP Storage Services Module        DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module              DS-X9112           ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
...
```

Step 5 Take the standby Supervisor-1 module out of service.

```
switch# out-of-service module 6
```

Step 6 Verify that the standby Supervisor-1 module is powered down.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module                DS-X9016           ok
2    32     Storage Services Module           DS-X9032-SSM       ok
3    8      IP Storage Services Module        DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module              DS-X9112           ok
5    0      Supervisor/Fabric-1                                  powered-dn
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
...
```

Step 7 Remove the standby Supervisor-1 module from the chassis.

Step 8 Install the Supervisor-2 module in the chassis.

Step 9 Establish a console session on the standby Supervisor-2 module console port (see the "Accessing the Switch" section on page 5-14).

Step 10 If the `loader>` prompt appears on the standby Supervisor-2 module console session, perform the following steps. Otherwise continue to Step 11.

a. Verify the Cisco SAN-OS system image and kickstart image are on the standby Supervisor-2 module bootflash:.

```
loader> dir
bootflash:

       14458880   kickstart-img
       12288      lost+found/
       48522513   system-img
```

b. If the images are present on bootflash:, boot the standby Supervisor-2 module using the **boot** command and continue to Step 11. Otherwise, continue to Step c.

```
loader> boot bootflash:kickstart-img bootflash:system-img
```

c. Enter the local IPv4 address, IPv4 subnet mask, and IPv4 address for the default gateway for the switch using the **network** command.

```
loader> network --ip=10.16.1.2 --nm=255.255.255.0 --gw=10.16.1.1
```

d. Boot the kickstart image file from the bootflash: (if present) or from a server.

```
loader> boot tftp://10.16.10.100/kickstart-img
INFO: Using IPv4 domain for downloading.
INFO: Initializing network interface...
INFO: Board will come up MGMT interface
INFO: MAC address is: 00:05:30:01:a8:0a
INFO: Ethernet port 1 link up and running at 100 Mbps full duplex.
TFTP from server 10.16.10.100; our IP address is 10.16.1.2; sending
through gateway 10.16.1.1
Filename 'kickstart-img'.
Load address: 0x400000

Loading: T ################T ##################################################
##################################################T ##################
###################T ##################################################
####################################################################
###################################T ##############################
####################################################################
#########################################################T ##############
####################################################################
#########################T ##########################################
########################################
done
Bytes transferred = 16061952 (f51600 hex)


### Network statistics: ###
--------------------------
Packets received: 31679
Packets send: 31389
Received bytes: 33167191
Send bytes: 1444202
Rx mulicast frames: 1
```

```
Automatic boot of image at addr 0x00000000 ...
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
WARNING: image sync is going to be disabled after a mgmt0 BIOS netboot
Loading system software
No system image is specified
INIT: Sending processes the TERM signal
Stopping kernel log daemon: klogd.
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2006, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
switch(boot)#
```

The `switch(boot)#` prompt indicates that you have a usable kickstart image.

e. Enable the management interface (mgmt0).

```
switch(boot)# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(boot)(config)# interface mgmt 0
switch(boot)(config-if)# no shutdown
switch(boot)(config-if)# end
switch(boot)#
```

f. Download a Cisco SAN-OS system image to the Supervisor-2 module from a TFTP server.

```
switch(boot)# copy tftp://10.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server......
```

Note Ensure that you download the correct Cisco SAN-OS system image. The system image for Supervisor-2 modules is different from the system image for Supervisor-1 modules.

g. Download a kickstart image to the Supervisor-2 module from a TFTP server, if necessary.

```
switch(boot)# copy tftp://10.16.10.100/kickstart-img bootflash:kickstart-img
Trying to connect to tftp server......
```

h. Load the system image on the standby Supervisor-2 module.

```
switch(boot)# load bootflash:system-img
```

i. Boot the standby Supervisor-2 module.

```
switch(boot)# boot bootflash:kickstart-img bootflash:system-img
```

Step 11 Verify that the standby Supervisor-2 module is in the warm standby state using a **show system redundancy status** command on the active Supervisor-1 module session.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative:   HA
```

```
        operational:   Warm

This supervisor (sup-2)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with warm standby

Other supervisor (sup-1)
------------------------
    Redundancy state:   Standby
    Supervisor state:   Warm standby
      Internal state:   Warm standby
```

Step 12 Copy the running configuration to the startup configuration on the active Supervisor-1 module to ensure that any running configuration changes are saved to the startup configuration and the ASCII configuration is synced and up-to-date on the warm standby Supervisor-2 module.

```
switch# copy running-config start-config
```

Step 13 If your switch has SSMs installed and intelligent services are configured, perform Step a through Step c. Otherwise, continue to Step 14.

a. Power down all SSMs on the switch.

```
switch# config t
switch(config)# poweroff module 2
switch(config)# exit
switch#
```

Caution Do not copy the running configuration to the startup configuration after powering down the SSMs. If you do, you will lose the configuration on the SSM interfaces.

b. Verify that the SSMs are powered down.

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     Storage Services Module                             powered-dn
3    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module             DS-X9112           ok
5    0      Supervisor/Fabric-2              DS-X9530-SF2-K9    ha-standby
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
...
```

c. Copy the contents of the SSM NVRAM to the standby Supervisor-2 module.

```
switch# copy ssm-nvram standby-sup
```

Step 14 Initiate a switchover on the active Supervisor-1 module to power it down and cause the standby Supervisor-2 module to become the active supervisor module.

```
switch# system switchover
```

Step 15 Verify that the Supervisor-1 module is powered down.

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    32     Storage Services Module          DS-X9032-SSM       ok
```

```
3    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    12     1/2/4 Gbps FC Module             DS-X9112           ok
5    0      Supervisor/Fabric-2              DS-X9530-SF2-K9    active *
6    0      Supervisor/Fabric-1                                 powered-dn
...
```

Step 16 Remove the Supervisor-1 module from the chassis.

Step 17 Set the baud rate on the active Supervisor-2 module console session to the default value of 9600.

```
switch# config t
switch(config)# line console
switch(config-console)# speed 9600
switch(config-console)# end
switch# show line console
line Console:
    Speed:        9600 bauds
    Databits:     8 bits per byte
    Stopbits:     1 bit(s)
    Parity:       none
    Modem In: Disable
    Modem Init-String -
        default : ATQ0V1H0S0=1\015
```

Step 18 Install the other Supervisor-2 module in the chassis.

Step 19 Verify that the standby Supervisor-2 module is in the HA-standby state.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative:   HA
         operational:   HA

This supervisor (sup-1)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with HA standby

Other supervisor (sup-2)
------------------------
    Redundancy state:   Standby
    Supervisor state:   HA standby
      Internal state:   HA standby
```

Step 20 If the Cisco MDS SAN-OS system image on the supervisor modules is the desired release, issue the **install all** command.

```
switch# install all
```

If you want a different release of the Cisco SAN-OS system image running on the switch, issue the **install all** command specifying the system image to perform a hitless upgrade (see the "Automated Upgrades" section on page 7-7).

```
switch# install all system tftp://10.16.10.100/system-img
```

Standby Supervisor Module Boot Variable Version

If the standby supervisor module boot variable images are not the *same* version as those running on the active supervisor module, the software forces the standby supervisor module to run the same version as the active supervisor module.

If you specifically set the boot variables of the standby supervisor module to a different version and reboot the standby supervisor module, the standby supervisor module will only load the specified boot variable if the same version is also running on the active supervisor module. At this point, the standby supervisor module is *not* running the images set in the boot variables.

Standby Supervisor Module Bootflash Memory

When updating software images on the standby supervisor module, verify that there is enough space available for the image using the **dir bootflash://sup-standby/** command. It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images. For information about displaying file systems and removing files, see the "Using Switch File Systems" section on page 2-27.

Standby Supervisor Module Boot Alert

If a standby supervisor module fails to boot, the active supervisor module detects that condition and generates a Call Home event and a system message and reboots the standby supervisor module approximately 3 to 6 minutes after the standby supervisor module moves to the `loader>` prompt.

The following system message is issued:

```
%DAEMON-2-SYSTEM_MSG:Standby supervisor failed to boot up.
```

This error message is also generated if one of the following situations apply:

- You remain at the `loader>` prompt for an extended period of time.
- You do not set the boot variables appropriately.

Installing Generation 2 Modules in Generation 1 Chassis

The Generation 2 modules have the following installation restrictions:

- Supervisor-2 modules can be installed on all Cisco MDS 9500 Series Directors.

Note The Cisco MDS 9513 Director does not support Supervisor-1 modules.

- Generation 2 switching modules can be installed on all Cisco MDS 9000 Family switches, except the Cisco MDS 9216 switch.
- Generation 1 modules can be used with Cisco MDS 9000 Family switches. However, installing Generation 1 modules in combination with Generation 2 switching modules in the same chassis reduces the capabilities of the Generation 2 switching modules (see the "About Combining Generation 1 and Generation 2 Switching Modules" section on page 14-16).
- Generation 1 and Generation 2 switching modules can be installed on Cisco MDS 9500 Family switches with either Supervisor-1 modules or Supervisor-2 modules.

Replacing Modules

When you replace any module (supervisor, switching, or services module), you must ensure that the new module is running the same software version as the rest of the switch.

Refer to *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for configuration details on replacing the caching services module (CSM).

Note When a spare standby supervisor module is inserted, it uses the same image as the active supervisor module. The Cisco SAN-OS software image is not automatically copied to the standby flash device.

Tip Issue the **install all** command to copy the Cisco SAN-OS software image to the standby supervisor module bootflash device.

Issuing the **install all** command after replacing any module, ensures the following actions:

- The proper system and kickstart images are copied on the standby bootflash: file system.
- The proper boot variables are set.
- The loader and the BIOS are upgraded to the same version available on the active supervisor module.

To replace a module in any switch in the Cisco MDS 9200 Series or 9500 Series, follow these steps:

Step 1 Create a backup of your existing configuration file, if required, using the **copy running-config startup-config** command.

Step 2 Replace the required module as specified in the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Step 3 Verify that space is available on the standby supervisor module bootflash using the **dir bootflash://sup-standby/** command. It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images. For information about displaying file systems and removing files, see the "Using Switch File Systems" section on page 2-27.

Step 4 Issue the **install all** command to ensure that the new module is running the same software as the rest of the switch.

Step 5 Wait until the new module is online and then ensure that the replacement was successful using the **show module** command.

Default Settings

Table 7-6 lists the default image settings for all Cisco MDS 9000 Family switches.

Table 7-6 Default Image Settings

Parameters	Default
Kickstart image	No image is specified.
System image	No image is specified.

CHAPTER 8

Working with Configuration Files

This chapter describes how to initially configure switches using the configuration files so they can be accessed by other devices. This chapter includes the following sections:

Managing Configuration Files

Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration so that they have identical module and port configurations.

This section describes how to work with configuration files and has the following topics:

Displaying Configuration Files

Use the **show running-config** command to view the running configuration file.

```
switch# show running-config
Building Configuration ...
interface port-channel 98
interface fc1/1
interface fc1/2
```

```
interface mgmt0
 ip address 172.22.95.112 255.255.255.0
 no shutdown
vsan database
 vsan 2
clock summer-time Pacific 1 Sun Apr 02:00 5 Sun Oct 02:00 60
switchname switch112
```

Use the **show startup-config** command to view the startup configuration file.

```
switch# show startup-config
interface port-channel 98
interface fc1/1
channel-group 98 force
 no shutdown
interface mgmt0
 ip address 172.22.95.112 255.255.255.0
boot system system-237; ep-41
boot kickstart boot-237 ep-41
ip domain-name cisco.com
```

Downloading Configuration Files to the Switch

You can configure a switch in the Cisco MDS 9000 Family by using configuration files you create or download from another switch. In addition, you can store configuration files on a bootflash device on the supervisor module and you can configure the switch using a configuration stored on an external CompactFlash disk.

Before you begin downloading a configuration file using a remote server, do the following:

- Ensure the configuration file to be downloaded is in the correct directory on the remote server.
- Ensure that the permissions on the file are set correctly. Permissions on the file should be set to world-read.
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router or default gateway to route traffic between subnets.

  Check connectivity to the remote server using the **ping** command.

From a Remote Server

To configure a switch in the Cisco MDS 9000 Family using a configuration file downloaded from a remote server using TFTP, FTP, SCP, or SFTP, follow these steps:

Step 1 Log into the switch through the console port or through a Telnet or SSH session.

Step 2 Configure the switch using the configuration file downloaded from the remote server using the **copy** *scheme***://***url* **system:running-config** command, where *scheme* is TFTP, FTP, SCP, or SFTP and *url* is the path to the source file on the remote server.

The configuration file downloads and the commands are executed as the file is parsed line by line.

Use the following command to download a configuration file from a remote server to the running configuration.

```
switch# copy scheme://url system:running-config
```

From an External CompactFlash Disk (slot0:)

Note The physical media must be inserted into slot0: after you log into the switch.

To configure a switch in the Cisco MDS 9000 Family using a configuration file stored on an external CompactFlash disk, follow these steps:

Step 1 Log into the switch through the console port or through a Telnet or SSH session.

Step 2 Insert the CompactFlash device into slot0: on the active supervisor module.

Step 3 Locate the configuration file using the **cd** and **dir** commands. (See the "Copying Configuration Files" section on page 8-5.)

Step 4 Configure the switch using the configuration file stored on the external CompactFlash disk using the **copy slot0:***source file* **system:running-config** command.

The commands are executed as the file is parsed line by line.

Use the following command to download a configuration file from an external CompactFlash to the running configuration:

```
switch copy slot0:dns-config.cfg system:running-config
```

Saving Configuration Files to an External Device

You can save a configuration file stored on internal storage to a remote server or to an external Flash device on the switch.

To a Remote Server

To save a configuration file to a remote server such as TFTP, FTP, SCP, or SFTP, follow these steps:

Step 1 Log into the switch through the console port or through a Telnet or SSH session.

Step 2 Save the configuration using the **copy system:running-config** *scheme***://***url* command, where *scheme* is TFTP, FTP, SCP, or SFTP and *url* is the path to the target file on the remote server.

Use the following command to save a running configuration file to a remote server:

```
switch# copy system:running-config scheme://url
```

Use the following command to save a startup configuration file to a remote server:

```
switch# copy nvram:startup-config scheme://url
```

To an External CompactFlash Disk (slot0:)

To save a configuration file on an external CompactFlash device, follow these steps:

Step 1 Log into the switch through the console port or through a Telnet session.

Step 2 Save the running-config file using the **copy system:running-config slot0:***destination file* command or the startup-config file using the **copy nvram:startup-config** *destination file*.

Use the following command to save a running configuration file to an external CompactFlash disk:

```
switch# copy system:running-config slot0:dns-config.cfg
```

Use the following command to save a startup configuration file to an external CompactFlash disk:

```
switch# copy nvram:startup-config slot0:dns-config.cfg
```

Saving the Running Configuration

After you have created a running configuration in system memory, you can save it to the startup configuration in NVRAM.

Use the following **copy** command to save the configuration to NVRAM:

```
switch# copy system:running-config nvram:startup-config
```

The **copy running-config startup-config** command is an alias to the previous command and is used frequently throughout this guide.

To cancel the copy operation initiated by another switch, use the following command:

```
switch# system startup-config abort
```

To cancel the operation locally and throughout the fabric, enter **Ctrl-c** on the console or Telnet session of the initiator switch.

See the "Preserving Module Configuration" section on page 11-7.

Saving Startup Configurations in the Fabric

You can use Cisco Fabric Services (CFS) to instruct the other switches in the fabric to save their configurations to their local NVRAM.

NVRAM using the following **copy** command:

```
switch# copy running-config startup-config fabric
```

Note If any remote switch in the fabric fails to complete the **copy running-config startup-config fabric** process, the request is discarded on the initiator switch and the failure errors are displayed in the initiator switch CLI session.

You can use the **show cfs application** command to verify that the Fabric Startup Configuration Manager (FSCM) application is enabled.

```
switch# show cfs application

----------------------------------------------
 Application    Enabled   Scope
----------------------------------------------
 ntp            No        Physical-all
 fscm           Yes       Physical-fc
 islb           No        Physical-fc
 role           No        Physical-all
 rscn           No        Logical
 radius         No        Physical-all
 fctimer        No        Physical-fc
 syslogd        No        Physical-all
 callhome       No        Physical-all
 fcdomain       No        Logical
 device-alias   Yes       Physical-fc

Total number of entries = 11
```

Unlocking the Startup Configuration File

The startup configuration file can be locked by applications on the switch. To display locks on the startup configuration file, use the following command:

```
switch# show system internal sysmgr startup-config locks
```

To release a lock on the startup configuration file, use the following command:

```
switch# system startup-config unlock 10
```

Copying Configuration Files

The syntax for the **copy** command follows and is explained in Table 8-1.

```
switch# copy scheme://server/filename scheme://server/filename
```

Table 8-1 copy Command Syntax

Scheme	Server	File Name
bootflash	sup-active sup-standby sup-1 or module-5 sup-2 or module-6 sup-local sup-remote	User-specified
slot0	—	User-specified
volatile	—	User-specified
nvram	—	startup-config or snapshot-config
system	—	running-config
tftp[1]	IPv4 address, IPv6 address, or DNS name	User-specified
ftp		
scp (secure copy)		
sftp		
core	*slot-number*	Process identifier number

1. When downloading and uploading files, a TFTP limitation restricts a TFTP client to a 32 MB file size and some TFTP servers to a 16-MB file size.

- This example shows how to copy a file from bootflash in the active supervisor module (sup-1 in slot 5 on Cisco MDS 9506 and Cisco MDS 9509 switches or slot 7 on Cisco MDS 9513 switches) to the bootflash in the standby supervisor module (sup-2 in slot 6 on Cisco MDS 9506 and Cisco MDS 9509 switches or slot 8 on Cisco MDS 9513 switches).

```
switch# copy bootflash:system_image bootflash://sup-2/system_image
```

- This example shows how to overwrite the contents of an existing configuration in NVRAM.

```
switch# copy nvram:snapshot-config nvram:startup-config
Warning: this command is going to overwrite your current startup-config.
Do you wish to continue? {y/n} [y] y
```

- This example shows how to copy a running configuration to the bootflash: file system.

```
switch# copy system:running-config bootflash:my-config
```

- This example shows how to copy a system image file from the SCP server identified by an IPv4 address to bootflash.

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```

- This example shows how to copy a script file from the SFTP server identified by an IPv4 address to the volatile: file system.

```
switch# copy sftp://172.16.10.100/myscript.txt volatile:myscript.txt
```

Note Use the **show version image** command to verify if the downloaded images are valid.

Backing UpConfiguration Files

All switch configurations reside in the internal system: and nvram: file systems. If your internal file systems are corrupted, you could potentially lose your configuration. Save and back up your configuration files periodically. Also, before installing or migrating to a new software configuration, back up the configuration files.

- This example shows how to create a snapshot of the startup configuration in a predefined location on the switch (binary file).

  ```
  switch# copy nvram:startup-config nvram:snapshot-config
  ```

- This example shows how to back up the startup configuration copy in the bootflash: file system (ASCII file).

  ```
  switch# copy nvram:startup-config bootflash:my-config
  ```

- This example shows how to back up the startup configuration to the TFTP server (ASCII file).

  ```
  switch# copy nvram:startup-config tftp://172.16.10.100/my-config
  ```

- This example shows how to back up the running configuration to the bootflash: file system (ASCII file).

  ```
  switch# copy system:running-config bootflash:my-config
  ```

Rolling Back to a Previous Configuration

Problems, such as memory corruption, can occur that make it necessary for you to recover your configuration from a backed up version.

This example shows how to roll back to a snapshot copy of a previously saved running configuration (binary file).

```
switch# copy nvram:snapshot-config nvram:startup-config
```

Note You can issue a rollback command only when a snapshot is already created. Otherwise, you will receive the `No snapshot-config found` error message.

Note Each time a **copy running-config startup-config** command is issued, a binary file is created and the ASCII file is updated. A valid binary configuration file reduces the overall boot time significantly. A binary file cannot be uploaded, but its contents can be used to overwrite the existing startup configuration. The **write erase** command clears the binary file.

Restoring the Configured Redundancy Mode

Tip If you configure the combined mode as the redundancy mode for power supplies on a Cisco MDS 9509 switch, be careful when using the **write erase** and **reload** command sequence before rolling back to a saved configuration because this command sequence causes the switch to revert to factory defaults.

By issuing the **write erase** command and the **reload** command, you restore the switch settings to their factory defaults. This sequence also restores the redundancy mode setting for the power supplies back to the redundant mode (default).

Depending on the type of power supply, the input voltage, and the number of modules (line cards) in the chassis, the redundancy mode may prevent the line cards from being powered on after a system reboot (see the "Power Supply Configuration Modes" section on page 10-11). If you use this sequence, the commands that apply to the powered down line cards will not be enforced on the switch (and will not be part of its running configuration).

If using the **write erase** and **reload** command sequence before rolling back to a saved configuration, follow these steps:

Step 1 Manually change (if originally configured) the redundant mode configuration to combined mode.

Step 2 Wait until all modules are back online—the module status displays `ok` in response to the **show module** command.

Step 3 Roll back to the saved configuration (see the "Rolling Back to a Previous Configuration" section on page 8-7).

Accessing File Systems on the Standby Supervisor Module

You can access all file systems on the standby supervisor module (remote) from a session on the active supervisor module. This is useful when copying files to the active supervisor modules requires similar files to exist on the standby supervisor module.

- Use the **dir** *scheme***://sup-remote** command to list files on the standby supervisor module.

```
switch# dir bootflash://sup-remote
   12198912     Aug 27 16:29:18 2003  m9500-sf1ek9-kickstart-mzg.1.3.0.39a.bin
    1864931     Apr 29 12:41:59 2003  dplug2
      12288     Apr 18 20:23:11 2003  lost+found/
   12097024     Nov 21 16:34:18 2003  m9500-sf1ek9-kickstart-mz.1.3.1.1.bin
   41574014     Nov 21 16:34:47 2003  m9500-sf1ek9-mz.1.3.1.1.bin
Usage for bootflash://sup-remote
   67747169 bytes used
  116812447 bytes free
  184559616 bytes total
```

- Use the **delete** *scheme***://sup-remote** to remove files from a file system on the standby supervisor module.

```
switch# delete bootflash://sup-remote/aOldConfig.txt
```

Deleting Configuration Files

Use the **delete** command to remove configuration files from the memory locations on the switch.

- This example shows how to delete a file from the bootflash: file system.

```
switch# delete bootflash:dns_config.cfg
```

- This example shows how to delete a file from an external CompactFlash (slot0:) file system.

```
switch# delete slot0:dns_config.cfg
```

- This example shows how to delete the file named test from the Flash card inserted in slot 0.

```
switch# delete slot0:test
Delete slot0:test? [y/n]: y
```

- This example shows how to delete the entire `my-dir` directory and all its contents.

```
switch# delete bootflash:my-dir
```

- This example shows how to delete a file in the bootflash: on the standby supervisor module.

```
switch# delete bootflash://sup-remote/aOldConfig
```

CHAPTER 9

Configuring High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework.

This chapter includes the following sections:

About High Availability

The high availability (HA) software framework provides the following:

- Ensures nondisruptive software upgrade capability. See Chapter 7, “Software Images.”
- Provides redundancy for supervisor module failure by using dual supervisor modules.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in switches in the Cisco MDS 9200 Series and the Cisco MDS 9100 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series. See Chapter 16, “Configuring PortChannels.”

- Provides management redundancy using the Virtual Router Redundancy Protocol (VRRP). This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series.

  See the "Virtual Router Redundancy Protocol" section on page 43-16.

- Provides switchovers if the active supervisor fails. The standby supervisor, if present, takes over without disrupting storage or host traffic.

  Directors in the Cisco MDS 9500 Series have two supervisor modules (sup-1 and sup-2) in slots 5 and 6 (Cisco MDS 9509 and 9506 Switches) or slots 7 and 8 (Cisco MDS 9513 Switch). When the switch powers up and both supervisor modules are present, the supervisor module that comes up first enters the active mode and the supervisor module that comes up second enters the standby mode. If both supervisor modules come up at the same time, sup-1 becomes active. The standby supervisor module constantly monitors the active supervisor module. If the active supervisor module fails, the standby supervisor module takes over without any impact to user traffic.

Note For high availability, you need to connect the ethernet port for both active and standby supervisors to the same network or virtual LAN. The active supervisor owns the one IP address used by these ethernet connections. On a switchover, the newly activated supervisor takes over this IP address.

Switchover Mechanisms

Switchovers occur by one of the following two mechanisms:

- The active supervisor module fails and the standby supervisor module automatically takes over.
- You manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a switchover process has started another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

Caution If the standby supervisor module is not in a stable state (ha-standby), a switchover is not performed.

HA Switchover Characteristics

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because control traffic is not impacted.
- It does not disrupt data traffic because the switching modules are not impacted.
- Switching modules are not reset.

Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, issue the **system switchover** command. Once issued, another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

To ensure that an HA switchover is possible, issue the **show system redundancy status** command or the **show module** command. If the command output displays the `HA-standby` state for the standby supervisor module, then the switchover is possible.

Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- When you manually initiate a switchover, system messages indicate the presence of two supervisor modules.
- A switchover can only be performed when two supervisor modules are functioning in the switch.
- The modules in the chassis are functioning as designed.

Verifying Switchover Possibilities

This section describes how to verify the status of the switch and the modules before a switchover.

- Use the s**how system redundancy status** command to ensure that the system is ready to accept a switchover.
- Use the **show module** command to verify the status (and presence) of a module at any time. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
2    8      IP Storage Services Module       DS-X9308-SMIP      ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
8    0      Caching Services Module          DS-X9560-SMAP      ok
9    32     1/2 Gbps FC Module               DS-X9032           ok

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
2    1.3(0.106a)  0.206   20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a)  0.602   --
6    1.3(0.106a)) 0.602   --
8    1.3(0.106a)  0.702   --
9    1.3(0.106a)  0.3     22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de  JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2  JAB06350B1R
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd  JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe  JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6  JAB06280ae9

* this terminal session
```

  The `Status` column in the output should display an `OK` status for switching modules and an active or `HA-standby` status for supervisor modules. If the status is either `OK` or active, you can continue with your configuration.

- Use the **show boot auto-copy** command to verify the configuration of the auto-copy feature and if an auto-copy to the standby supervisor module is in progress. Sample outputs of the **show boot auto-copy** command follow:

```
switch# show boot auto-copy
Auto-copy feature is enabled
switch# show boot auto-copy list
No file currently being auto-copied
```

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This vital process functions on infrastructure that is internal to the switch.

See the "Displaying System Processes" section on page 59-1.

Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module automatically synchronizes its image with the running image on the active supervisor module.

See the "Replacing Modules" section on page 7-41.

Copying Boot Variable Images to the Standby Supervisor Module

You can copy the boot variable images that are in the active supervisor module (but not in the standby supervisor module) to the standby supervisor module. Only those KICKSTART and SYSTEM boot variables that are set for the standby supervisor module can be copied. For module (line card) images, all boot variables are copied to the corresponding standby locations (bootflash: or slot0:) if not already present.

Automatic Copying of Boot Variables

To enable or disable automatic copying of boot variables, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# boot auto-copy` `Auto-copy administratively enabled`	Enables (default) automatic copying of boot variables from the active supervisor module to the standby supervisor module.
	`switch(config)# no boot auto-copy` `Auto-copy administratively disabled`	Disables the automatic copy feature.

Verifying the Copied Boot Variables

Use the **show boot auto-copy** command to verify the current state of the copied boot variables. This example output shows that automatic copying is enabled.

```
switch# show boot auto-copy
Auto-copy feature enabled
```

This example output shows that automatic copying is disabled.

```
switch# show boot auto-copy
Auto-copy feature disabled
```

Use the **show boot auto-copy list** command to verify what files are being copied. This example output displays the image being copied to the standby supervisor module's bootflash. Once this is successful, the next file will be image2.bin.

Note This command only displays files on the active supervisor module.

```
switch# show boot auto-copy list
File: /bootflash:/image1.bin
Bootvar: kickstart

File:/bootflash:/image2.bin
Bootvar: system
```

This example output displays a typical message when the **auto-copy** option is disabled or if no files are copied.

```
switch# show boot auto-copy list
No file currently being auto-copied
```

Displaying HA Status Information

Use the **show system redundancy status** command to view the HA status of the system. Tables 9-1 to 9-3 explain the possible output values for the redundancy, supervisor, and internal states.

```
switch# show system redundancy status
Redundancy mode
---------------
      administrative:   HA
         operational:   HA
This supervisor (sup-1)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with HA standby
Other supervisor (sup-2)
------------------------
    Redundancy state:   Standby
    Supervisor state:   HA standby
      Internal state:   HA standby
```

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is `Active with HA standby` and the other supervisor module is `HA-standby`, the switch is operationally HA and can do automatic synchronization.

- If the internal state of one of the supervisor modules is `none`, the switch cannot do automatic synchronization.

Table 9-1 lists the possible values for the redundancy states.

Table 9-1 Redundancy States

State	Description
Not present	The supervisor module is not present or is not plugged into the chassis.
Initializing	The diagnostics have passed and the configuration is being downloaded.
Active	The active supervisor module and the switch is ready to be configured.
Standby	A switchover is possible.
Failed	The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state.
Offline	The supervisor module is intentionally shut down for debugging purposes.
At BIOS	The switch has established connection with the supervisor and the supervisor module is performing diagnostics.
Unknown	The switch is in an invalid state. If it persists, call TAC.

Table 9-2 lists the possible values for the supervisor module states.

Table 9-2 Supervisor States

State	Description
Active	The active supervisor module in the switch is ready to be configured.
HA standby	A switchover is possible.
Offline	The switch is intentionally shut down for debugging purposes.
Unknown	The switch is in an invalid state and requires a support call to TAC.

Table 9-3 lists the possible values for the internal redundancy states.

Table 9-3 Internal States

State	Description
HA standby	The HA switchover mechanism in the standby supervisor module is enabled (see the "HA Switchover Characteristics" section on page 9-2).
Active with no standby	A switchover is possible.
Active with HA standby	The active supervisor module in the switch is ready to be configured. The standby supervisor module is in the HA-standby state.
Shutting down	The switch is being shut down.
HA switchover in progress	The switch is in the process of changing over to the HA switchover mechanism.
Offline	The switch is intentionally shut down for debugging purposes.

Table 9-3 Internal States (continued)

State	Description
HA synchronization in progress	The standby supervisor module is in the process of synchronizing its state with the active supervisor modules.
Standby (failed)	The standby supervisor module is not functioning.
Active with failed standby	The active supervisor module and the second supervisor module is present but is not functioning.
Other	The switch is in a transient state. If it persists, call TAC.

CHAPTER 10

Managing System Hardware

This chapter provides details on how to manage system hardware other than services and switching modules and how to monitor the health of the switch. It includes the following sections:

Displaying Switch Hardware Inventory

Use the **show inventory** command to view information on the field replaceable units (FRUs) in the switch, including product IDs, serial numbers, and version IDs. See Example 10-1.

Example 10-1 Displays the Hardware Inventory

```
switch# show inventory
NAME: "Chassis",  DESCR: "MDS 9506 chassis"
PID: DS-C9506            ,  VID: 0.104,  SN: FOX0712S00T

NAME: "Slot 3",  DESCR: "2x1GE IPS, 14x1/2Gbps FC Module"
PID: DS-X9302-14K9       ,  VID: 0.201,  SN: JAB081405AF

NAME: "Slot 4",  DESCR: "2x1GE IPS, 14x1/2Gbps FC Module"
PID: DS-X9302-14K9       ,  VID: 0.201,  SN: JAB081605A5

NAME: "Slot 5",  DESCR: "Supervisor/Fabric-1"
PID: DS-X9530-SF1-K9     ,  VID: 4.0,  SN: JAB0747080H
```

```
NAME: "Slot 6",  DESCR: "Supervisor/Fabric-1"
PID: DS-X9530-SF1-K9     ,  VID: 4.0,  SN: JAB0746090H

NAME: "Slot 17",  DESCR: "MDS 9506 Power Supply"
PID: DS-CAC-1900W        ,  VID: 1.0,  SN: DCA07216052

NAME: "Slot 19",  DESCR: "MDS 9506 Fan Module"
PID: DS-6SLOT-FAN        ,  VID: 0.0,  SN: FOX0638S150
```

Use the **show hardware** command to display switch hardware inventory details. See Example 10-2.

Note To display and configure modules, see Chapter 11, "Managing Modules."

Example 10-2 Displays the Hardware Information

```
switch# show hardware
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2003-2004 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license.


Software
  BIOS:      version 1.0.8
  loader:    version 1.1(0.114)
  kickstart: version 1.3(4a)
  system:    version 1.3(4a)

  BIOS compile time:       08/07/03
  kickstart image file is: bootflash:///boot-17r
  kickstart compile time:  10/25/2010 12:00:00
  system image file is:    bootflash:///isan-17r
  system compile time:     10/25/2020 12:00:00


Hardware
  RAM 1024592 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:           0 blocks (block size 512b)

  172.22.90.21 uptime is 7 days 4 hours 48 minute(s) 2 second(s)

  Last reset at 272247 usecs after Thu Sep 11 21:47:05 1980
    Reason: Reset Requested by CLI command reload
    System version: 1.3(4a)


This supervisor carries Pentium processor with 1024592 kB of memory
Intel(R) Pentium(R) III CPU at family with 512 KB L2 Cache
Rev: Family  6, Model  11 stepping  1

512K bytes of non-volatile memory.
1000944 blocks of internal bootflash (block size 512b)

--------------------------------
Chassis has 9 slots for Modules
--------------------------------
```

```
Module in slot 1 is empty

Module in slot 2 is empty

Module in slot 3 is empty

Module in slot 4 is empty

Module in slot 5 is ok
  Module type is "Supervisor/Fabric-1"
  No submodules are present
  Model number is DS-X9530-SF1-K9
  H/W version is 1.0
  Part Number is 73-7523-06
  Part Revision is A0
  Manufacture Date is Year 6 Week 47
  Serial number is JAB064705E1
  CLEI code is CNP6NT0AAA

Module in slot 6 is empty

Module in slot 7 is empty

Module in slot 8 is empty

Module in slot 9 is empty

--------------------------------------
Chassis has 2 Slots for Power Supplies
--------------------------------------

PS in slot A is ok
  Power supply type is "1153.32W 110v AC"
  Model number is WS-CAC-2500W
  H/W version is 1.0
  Part Number is 34-1535-01
  Part Revision is A0
  Manufacture Date is Year 6 Week 16
  Serial number is ART061600US
  CLEI code is

PS in slot B is ok
  Power supply type is "1153.32W 110v AC"
  Model number is WS-CAC-2500W
  H/W version is 1.0
  Part Number is 34-1535-01
  Part Revision is A0
  Manufacture Date is Year 5 Week 41
  Serial number is ART0541003V
  CLEI code is

---------------------------------
Chassis has one slot for Fan Module
---------------------------------

  Fan module is ok
  Model number is WS-9SLOT-FAN
  H/W version is 0.0
  Part Number is 800-22342-01
  Part Revision is
  Manufacture Date is Year 0 Week 0
  Serial number is
  CLEI code is
```

Running Compact Flash Tests

In Cisco SAN-OS Release 3.1(3), you can run the CompactFlash CRC checksum test to identify if the CompactFlash firmware is corrupted and needs to be updated. By default, the CompactFlash CRC checksum test is enabled to automatically run in the background every seven days (you can change the automatic test interval by using the **system health module cf-crc-check frequency** command in configuration mode). You can run the test on demand by using the **system health cf-crc-check module** CLI command in EXEC mode. To turn the automatic testing off, use the **no system health module cf-crc-check** command in configuration mode.

The CompactFlash CRC checksum test can check if CompactFlash is corrupted on the following modules:

- DS-X9016
- DS-X9032
- DS-X9302-14K9
- DS-X9308-SMIP
- DS-X9304-SMIP
- DS-X9530-SF1-K9

This section includes the following tasks:

Running the CompactFlash CRC Checksum Test On Demand

To run the CompactFlash CRC checksum test, use the **system health cf-crc-check module** command in EXEC mode.

```
switch# system health cf-crc-check module number
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch(config)# system health cf-crc-check module 4
```

Enabling and Disabling the Automatic CompactFlash CRC Checksum Test

By default, the CompactFlash CRC Checksum test is enabled to automatically run in the background. You can disable the automatic testing and then enable the testing at a later time. You can run the test on demand at any time by using the **system health cf-crc-check module** command in EXEC mode.

To enable automatic CompactFlash CRC checksum testing, use the **system health module cf-crc-check** command in configuration mode.

```
switch# system health module number cf-crc-check
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch(config)# system health module 4 cf-crc-check
```

To disable automatic CompactFlash CRC checksum testing, use the **no system health module cf-crc-check** command in EXEC mode.

```
switch(config)# no system health module number cf-crc-check
```

Setting the CompactFlash CRC Checksum Test Interval

To set the CompactFlash CRC checksum test interval, use the **system health module cf-crc-check frequency** command in configuration mode.

To set the CompactFlash CRC checksum test interval, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system health module number cf-crc-check frequency number` `switch(config)#`	Configures the CompactFlash CRC checksum test interval in days. The default interval is 7 days.

Enabling and Disabling Failure Action at the Failure of a CompactFlash Checksum Test

You can use the **system health module cf-crc-check failure-action** command to prevent the Cisco SAN-OS software from taking any action if a CompactFlash failure is determined while running the CRC checksum test and the failed CompactFlash is isolated from further testing. By default, this feature is enabled in all switches in the Cisco MDS 9000 Family. A failure action is controlled at the module level.

Use the **system health module cf-crc-check failure-action** command in configuration mode to enable the CompactFlash CRC checksum test failure action for a module.

To enable the CompactFlash CRC checksum test failure action, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system health module number cf-crc-check failure-action` `switch(config)#`	Enables the CompactFlash CRC checksum test failure action for a specified module.

To disable the CompactFlash CRC checksum test failure action, use the **no system health module cf-crc-check failure-action** command in configuration mode.

Displaying the Frequency and Status of the CompactFlash CRC Checksum Test

To display the frequency and status of the CompactFlash CRC checksum test for a specific module, use the **show system health module** command in EXEC mode.

```
switch# show system health module 5

Current health information for module 5.
```

```
Test                    Frequency       Status          Action
------------------------------------------------------------------
Bootflash                10 Sec         Running         Enabled
EOBC                      5 Sec         Running         Enabled
Loopback                  5 Sec         Running         Enabled
CF checksum               1 Day         Running         Enabled
CF re-flash              30 Days        Running         Enabled
------------------------------------------------------------------
```

Updating the CompactFlash Firmware

In Cisco SAN-OS Release 3.1(3), you can update the CompactFlash firmware on selected modules. By default, the firmware update feature is enabled to automatically update the firmware every 30 days (you can manually set the firmware update intervals by using the **system health module cf-re-flash frequency** command in configuration mode). You can also update the firmware on demand by using the **system health cf-re-flash module** command in EXEC mode.

Firmware updates can be enabled on the following modules:

- DS-X9016
- DS-X9032
- DS-X9302-14K9
- DS-X9308-SMIP
- DS-X9304-SMIP
- DS-X9530-SF1-K9

This section includes the following tasks:

Updating the CompactFlash Firmware On Demand

To update the CompactFlash firmware on demand, use the **system health cf-re-flash module** command in EXEC mode.

```
switch# system health cf-re-flash module number
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch(config)# system health cf-re-flash module 4
```

Enabling and Disabling the CompactFlash Firmware Update

By default, the CompactFlash firmware is updated automatically every30 days. You can disable the automatic update and then enable the automatic update at a later time. You can update the CompactFlash firmware on demand at any time by using the **system health cf-re-flash module** command in EXEC mode.

To enable automatic firmware updates, use the **system health module cf-re-flash** command in configuration mode.

```
switch(config)# system health module number cf-re-flash
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch(config)# system health module 4 cf-re-flash
```

To disable automatic firmware updates, use the **no system health module cf-re-flash** command in configuration mode.

```
switch(config)# no system health module number cf-re-flash
```

Setting the CompactFlash Firmware Update Interval

To set the firmware update interval, use the **system health module cf-re-flash frequency** command in configuration mode. The default interval is every 30 days.

To set the firmware update interval, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system health module number cf-re-flash frequency number` `switch(config)#`	Configures the firmware update interval in days. The default interval is 30 days.

Enabling and Disabling Failure Action at the Failure of a CompactFlash Firmware Update

You can use the **system health module cf-re-flash failure-action** command to prevent the Cisco SAN-OS software from taking any action if a CompactFlash failure is determined while updating the CompactFlash firmware. By default, this action is taken if a failure is determined and the failed CompactFlash is isolated from further testing. A failure action is controlled at the module level.

Use the **system health module cf-re-flash failure-action** command in configuration mode to enable the CompactFlash firmware update failure action for a module.

To enable the CompactFlash CRC checksum test failure action, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system health module number cf-re-flash failure-action` `switch(config)#`	Enables the CompactFlash CRC firmware update failure action for a specified module.

To disable the CompactFlash CRC firmware update test failure action, use the **no system health module cf-re-flash failure-action** command in configuration mode.

Displaying the Frequency and Status of CompactFlash Updates

To display the frequency and status of the flash updates for a specific module, use the **show system health module** command in EXEC mode.

```
switch# show system health module 5

Current health information for module 5.

Test                    Frequency       Status          Action
------------------------------------------------------------------
Bootflash                10 Sec         Running         Enabled
EOBC                      5 Sec         Running         Enabled
Loopback                  5 Sec         Running         Enabled
CF checksum               1 Day         Running         Enabled
CF re-flash              30 Days        Running         Enabled
------------------------------------------------------------------
```

Displaying CompactFlash CRC Test and Firmware Update Statistics

To display the CompactFlash CRC checksum test and the flash update statistics, use the **show system health statistics** command in EXEC mode.

```
switch# show system health statistics

Test statistics for module 2
-------------------------------------------------------------------------------
Test Name           State            Freqency   Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
Bootflash           Running            10s    28316   28316       0     0    0
EOBC                Running             5s    56632   56632       0     0    0
Loopback            Running             5s    56618   56618       0     0    0
CF checksum         Running             2d        2       2       0     0    0
CF re-flash         Running            30d        1       1       0     0    0
-------------------------------------------------------------------------------

Test statistics for module 5
-------------------------------------------------------------------------------
Test Name           State            Freqency   Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
Bootflash           Running            10s    28314   28314       0     0    0
EOBC                Running             5s    56629   56629       0     0    0
Loopback            Running             5s    56614   56614       0     0    0
CF checksum         Running             1d        4       4       0     0    0
CF re-flash         Running            30d        1       1       0     0    0
-------------------------------------------------------------------------------

Test statistics for module 7
-------------------------------------------------------------------------------
Test Name           State            Freqency   Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
InBand              Running             5s    56643   56643       0     0    0
Bootflash           Running            10s    28323   28323       0     0    0
EOBC                Running             5s    56643   56643       0     0    0
Management Port     Running             5s    56643   56643       0     0    0
-------------------------------------------------------------------------------
```

```
Test statistics for module 8
-----------------------------------------------------------------------------
Test Name           State            Freqency   Run    Pass    Fail CFail Errs
-----------------------------------------------------------------------------
InBand              Running             5s    56624   56624       0     0    0
Bootflash           Running            10s    28317   28317       0     0    0
EOBC                Running             5s    56624   56624       0     0    0
-----------------------------------------------------------------------------

Test statistics for module 13
-----------------------------------------------------------------------------
Test Name           State            Freqency   Run    Pass    Fail CFail Errs
-----------------------------------------------------------------------------
Bootflash           Running            10s    28304   28304       0     0    0
EOBC                Running             5s    56608   56608       0     0    0
Loopback            Running             5s    56608   56608       0     0    0
-----------------------------------------------------------------------------
```

Displaying the Switch Serial Number

The serial number of your Cisco MDS 9000 Family switch can be obtained by looking at the serial number label on the back of the switch (next to the power supply), or by executing the operating system **show sprom backplane 1** command.

```
switch# show sprom backplane 1
DISPLAY backplane sprom contents:
Common block:
 Block Signature : 0xabab
 Block Version   : 2
 Block Length    : 156
 Block Checksum  : 0x106f
 EEPROM Size     : 512
 Block Count     : 3
 FRU Major Type  : 0x6001
 FRU Minor Type  : 0x0
 OEM String      : Cisco Systems, Inc.
 Product Number  : DS-C9506
 Serial Number   : FOX0712S007
 Part Number     : 73-8697-01
 Part Revision   : 01
 Mfg Deviation   : 0
 H/W Version     : 0.1
 Mfg Bits        : 0
 Engineer Use    : 0
 snmpOID         : 9.12.3.1.4.26.0.0
 Power Consump   : 0
 RMA Code        : 0-0-0-0
Chassis specific block:
...
```

Note If you are installing a new license, use the **show license host-id** command to obtain the switch serial. See Chapter 3, “Obtaining and Installing Licenses,” for further information.

Displaying Power Usage Information

Use the **show environment power** command to display the actual power usage information for the entire switch. In response to this command, power supply capacity and consumption information is displayed for each module. See Example 10-3.

Note In a Cisco MDS 9500 Series switch, power usage is reserved for both supervisors regardless of whether one or both supervisor modules are present.

Example 10-3 Displays Power Management Information

```
switch# show environment power

-------------------------------------------------
PS  Model                Power     Power       Status
                         (Watts)   (Amp @42V)
-------------------------------------------------
1   DS-CAC-2500W         1153.32   27.46       ok
2   WS-CAC-2500W         1153.32   27.46       ok


Mod Model                Power     Power       Power     Power       Status
                         Requested Requested   Allocated Allocated
                         (Watts)   (Amp @42V)  (Watts)   (Amp @42V)
--- -------------------  -------   ----------  --------- ----------  ----------
1    DS-X9032            199.92    4.76        199.92    4.76        powered-up
4    DS-X9032            199.92    4.76        199.92    4.76        powered-up
5    DS-X9530-SF1-K9     126.00    3.00        126.00    3.00        powered-up
6    DS-X9530-SF1-K9     126.00    3.00        126.00    3.00        powered-up
9    DS-X9016            220.08    5.24        220.08    5.24        powered-up


Power Usage Summary:
--------------------
Power Supply redundancy mode:                  redundant

Total Power Capacity                           1153.32  W

Power reserved for Supervisor(s)[-]            252.00  W
Power reserved for Fan Module(s) [-]           0.00    W
Power currently used by Modules[-]             619.92  W

                                               -------
Total Power Available                          281.40   W
                                               -------
```

Power Supply Configuration Modes

Switches in the MDS 9000 Family have two redundant power supply slots. The power supplies can be configured in either redundant or combined mode.

- Redundant mode—Uses the capacity of one power supply only. This is the default mode. In case of power supply failure, the entire switch has sufficient power available in the system.
- Combined mode—Uses the combined capacity of both power supplies. In case of power supply failure, the entire switch can be shut down (depends on the power used) causing traffic disruption. This mode is seldom used, except in cases where the switch has two low power supply capacities but a higher power usage.

Note The chassis in the Cisco MDS 9000 Family uses 1200 W when powered at 110 V, and 2500 W when powered at 220 V.

To configure the power supply mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# power redundancy-mode combined` `switch(config)#`	Configures combined power supply mode.
	`switch(config)# power redundancy-mode redundant` `switch(config)#`	Reverts to the redundant (default) power supply mode.

Note See the Use the **show environment power** command to view the current power supply configuration.

Power Supply Configuration Guidelines

Follow these guidelines when configuring power supplies:

1. When power supplies with different capacities are installed in the switch, the total power available differs based on the configured mode, either redundant or combined:
   a. Redundant mode—the total power is the lesser of the two power supply capacities.
   For example, suppose you have the following usage figures configured:

   Power supply 1 = 2500 W
   Additional power supply 2 = not used
   Current usage = 2000 W
   Current capacity = 2500 W

   Then the following three scenarios differ as specified (see Table 10-1):

   Scenario 1: If 1800 W is added as power supply 2, then power supply 2 is shut down.
   Reason: 1800 W is less than the usage of 2000 W.

   Scenario 2: If 2200 W is added as power supply 2, then the current capacity decreases to 2200 W.
   Reason: 2200 W is the lesser of the two power supplies.

Scenario 3: If 3000 W is added as power supply 2, then the current capacity value remains at 2500 W.
Reason: 2500 W is the lesser of the two power supplies.

Table 10-1 Redundant Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Usage (W)	Insertion of Power Supply 2 (W)	New Capacity (W)	Action Taken by Switch
1	2500	2000	1800	2500	Power supply 2 is shut down.
2	2500	2000	2200	2200	Capacity becomes 2200 W.
3	2500	2000	3300	2500	Capacity remains the same.

1. W = Watts

b. Combined mode—the total power is twice the lesser of the two power supply capacities.

For example, suppose you have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = not used
Current Usage = 2000 W
Current capacity = 2500 W

Then, the following three scenarios differ as specified (see Table 10-2):

Scenario 1: If 1800 W is added as power supply 2, then the capacity increases to 3600 W.
Reason: 3600 W is twice the minimum (1800 W).

Scenario 2: If 2200 W is added as power supply 2, then the current capacity increases to 4400 W.
Reason:4400 W is twice the minimum (2200 W).

Scenario 3: If 3000 W is added as power supply 2, then the current capacity increases to 5000 W.
Reason: 5000 W is twice the minimum (2500 W).

Table 10-2 Combined Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Usage (W)	Insertion of Power Supply 2 (W)	New Capacity (W)	Action Taken by Switch
1	2500	2000	1800	3600	Power is never shut down. The new capacity is changed.
2	2500	2000	2200	4400	
3	2500	2000	3300	5000	

1. W = Watts

2. When you change the configuration from combined to redundant mode and the system detects a power supply that has a capacity lower than the current usage, the power supply is shut down. If both power supplies have a lower capacity than the current system usage, the configuration is not allowed. Several configuration scenarios are summarized in Table 10-3.

Scenario 1: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 1800 W
Current Usage = 2000 W
Current mode = combined mode (so current capacity is 3600 W)

You decide to change the switch to redundant mode. Then power supply 2 is shut down.

Reason: 1800 W is the lesser of the two power supplies and it is less than the system usage.

Scenario 2: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 2200 W
Current Usage = 2000 W
Current mode = combined mode (so current capacity is 4400 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2200 W.

Reason: 2200 W is the lesser of the two power supplies.

Scenario 3: You have the following usage figures configured:

Power supply 1 = 2500 W
Additional Power supply 2 = 1800 W
Current Usage = 3000 W
Current mode = combined mode (so current capacity is 3600 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2500 W and the configuration is rejected.

Reason: 2500 W is less than the system usage (3000 W).

Table 10-3 Combined Mode Power Supply Scenarios

Scenario	Power Supply 1 (W)[1]	Current Mode	Current Usage (W)	Power Supply 2 (W)	New Mode	New Capacity (W)	Action Taken by Switch
1	2500	combined	2000	1800	N/A	3600	This is the existing configuration.
	2500	N/A	2000	1800	redundant	2500	Power supply 2 is shut down
2	2500	combined	2000	2200	N/A	4400	This is the existing configuration.
	2500	N/A	2000	2200	redundant	2200	The new capacity is changed.
3	2500	combined	3000	1800	N/A	3600	This is the existing configuration.
	2500	N/A	3000	1800	redundant	N/A	Rejected, so the mode reverts to combined mode.

1. W = Watts

About Crossbar Management

Cisco MDS SAN-OS Release 3.0(1) and later supports two types of hardware for the Cisco MDS 9500 Series Directors: Generation 1 and Generation 2.

Generation 1 consists of all hardware supported by Cisco SAN-OS prior to Release 3.0(1), including the following:

- Cisco MDS 9506 and 9509 Director chassis
- Supervisor-1 module
- 32-port 2-Gbps Fibre Channel switching module
- 16-port 2-Gbps Fibre Channel switching module
- 8-port IP Storage Services (IPS-8) module
- 4-port IP Storage Services (IPS-4) module
- Storage Services Module (SSM)
- 14/2-port Multiprotocol Services (MPS-14/2) module

Generation 2 consists of all new hardware supported by Cisco SAN-OS Release 3.0(1) and later, including the following:

- Cisco MDS 9513 Director chassis
- Supervisor-2 module
- 48-port 4-Gbps Fibre Channel switching module
- 24-port 4-Gbps Fibre Channel switching module
- 12-port 4-Gbps Fibre Channel switching module
- 4-port 10-Gbps Fibre Channel switching module

The Cisco MDS 9500 Series Directors running Cisco MDS SAN-OS 3.0(1) or later support the following types of crossbars:

- Integrated crossbar—Located on the Supervisor-1 and Supervisor-2 modules. The Cisco MDS 9506 and 9509 Directors only use integrated crossbars.
- External crossbar—Located on an external crossbar switching module. Cisco MDS 9513 Directors require external crossbar modules.

Operational Considerations When Removing Crossbars

You can mix and match Generation 1 and Generation 2 hardware on the Cisco MDS 9500 Series Directors running Cisco MDS SAN-OS 3.0(1) or later without compromising the integrity and availability of your SANs based on Cisco MDS 9500 Series Directors.

To realize these benefits, you must consider the following operational requirements when *removing* crossbars for maintenance activities:

- Graceful Shutdown of a Crossbar, page 10-15
- Backward Compatibility for Generation 1 Modules in Cisco MDS 9513 Directors, page 10-15

Graceful Shutdown of a Crossbar

You must perform a graceful shutdown of a crossbar (integrated or external) before removing it from the MDS 9500 Series Director.

- You must enter the EXEC mode **out-of-service xbar** command for a graceful shutdown of external crossbar modules in a Cisco MDS 9513 Director.

  out-of-service xbar *slot*

  Where *slot* indicates the external crossbar module slot number.

  Note To reactivate the external crossbar module, you must remove and reinsert or replace the crossbar module.

- You must enter the EXEC mode **out-of-service module** command for a graceful shutdown of integrated crossbars on the supervisor module in a Cisco MDS 9506 or 9509 Director.

  out-of-service module *slot*

  Where *slot* indicates the chassis slot number on either the Supervisor-1 module or the Supervisor-2 module in which the integrated crossbar resides.

  Note To reactivate the integrated crossbar, you must remove and reinsert or replace the Supervisor-1 module or Supervisor-2 module.

Caution Taking the crossbar out-of-service may cause a supervisor switchover.

Backward Compatibility for Generation 1 Modules in Cisco MDS 9513 Directors

To provide backward compatibility for a Generation 1 module in a Cisco MDS 9513 chassis, the active and backup Supervisor-2 modules are associated to a specific crossbar module. The Supervisor-2 module in slot 7 is associated with crossbar module 1 and Supervisor-2 module in slot 8 is associated with crossbar module 2. You must plan for the following operational considerations before removing crossbar modules:

- Whenever a crossbar module associated with the active Supervisor-2 module goes offline or is brought online in a system that is already online, a stateful supervisor switchover occurs. This switchover does not disrupt traffic. Events that cause a crossbar module to go offline include the following:
  - Out-of-service requests
  - Physical removal
  - Errors
- Supervisor-2 module switchovers do *not* occur if the crossbar switching module associated with the *backup* Supervisor-2 module goes offline.

Note Supervisor-2 module switchovers do *not* occur when removing crossbar switch modules on a Cisco MDS 9513 that only has Generation 2 modules installed.

About Module Temperature

Built-in, automatic sensors are provided in all switches in the Cisco MDS 9000 Family to monitor your switch at all times.

Each module (switching and supervisor) has four sensors: 1 (outlet sensor), 2 (intake sensor), 3 (onboard sensor), and 4 (onboard sensor). Each sensor has two thresholds (in degrees Celsius): minor and major.

Note A threshold value of -127 indicates that no thresholds are configured or applicable.

- Minor threshold—When a minor threshold is exceeded, a minor alarm occurs and the following action is taken for all four sensors:
  - System messages are displayed.
  - Call Home alerts are sent (if configured).
  - SNMP notifications are sent (if configured).
- Major threshold—When a major threshold is exceeded, a major alarm occurs and the following action is taken.
  - For sensors 1, 3, and 4 (outlet and onboard sensors):

    System messages are displayed.

    Call Home alerts are sent (if configured).

    SNMP notifications are sent (if configured).
  - For sensor 2 (intake sensor):

    If the threshold is exceeded in a switching module, only that module is shut down.

    If the threshold is exceeded in an active supervisor module with HA-standby or standby present, only that supervisor module is shut down and the standby supervisor module takes over.

    If you do not have a standby supervisor module in your switch, you have an interval of 2 minutes to decrease the temperature. During this interval the software monitors the temperature every five (5) seconds and continuously sends system messages as configured.

Tip To realize the benefits of these built-in, automatic sensors on any switch in the Cisco MDS 9500 Series, we highly recommend that you install dual supervisor modules. If you are using a Cisco MDS 9000 Family switch without dual supervisor modules, we recommend that you immediately replace the fan module if even one fan is not working.

Displaying Module Temperature

Use the **show environment temperature** command to display temperature sensors for each module (see Example 10-4 and Example 10-5).

Example 10-4 Displays Temperature Information for Generation 1 Hardware

```
switch# show environment temperature
-----------------------------------------------------------------
Module   Sensor   MajorThresh   MinorThres   CurTemp     Status
                  (Celsius)     (Celsius)    (Celsius)
-----------------------------------------------------------------
2        Outlet   75            60           35          ok
2        Intake   65            50           33          ok

5        Outlet   75            60           44          ok
5        Intake   65            50           36          ok

6        Outlet   75            60           42          ok
6        Intake   65            50           35          ok

7        Outlet   75            60           33          ok
7        Intake   65            50           30          ok

9        Outlet   75            60           34          ok
9        Intake   65            50           39          ok
```

Example 10-5 Displays Temperature Information for Generation 2 Hardware

```
switch# show environment temperature
-----------------------------------------------------------------
Module   Sensor   MajorThresh   MinorThres   CurTemp     Status
                  (Celsius)     (Celsius)    (Celsius)
-----------------------------------------------------------------
1        Outlet1  75            60           33          ok
1        Outlet2  65            50           30          ok
1        Intake1  65            50           30          ok
1        LcFwdUp  65            50           35          ok
1        LcFwdDn  65            50           39          ok
1        FC-MAC   65            50           34          ok

6        Outlet1  75            60           33          ok
6        Outlet2  65            50           30          ok
6        Intake1  65            50           30          ok
6        Crosbar  65            50           35          ok
6        Arbiter  65            50           39          ok
6        CPU      65            50           34          ok
```

About Fan Modules

Hot-swappable fan modules (fan trays) are provided in all switches in the Cisco MDS 9000 Family to manage airflow and cooling for the entire switch. Each fan module contains multiple fans to provide redundancy. The switch can continue functioning in the following situations:

- One or more fans fail within a fan module—Even with multiple fan failures, switches in the Cisco MDS 9000 Family can continue functioning. When a fan fails within a module, the functioning fans in the module increase their speed to compensate for the failed fan(s).
- The fan module is removed for replacement—The fan module is designed to be removed and replaced while the system is operating without presenting an electrical hazard or damage to the system. When replacing a failed fan module in a running switch, be sure to replace the new fan module within five minutes.

Tip If one or more fans fail within a fan module, the Fan Status LED turns red. A fan failure could lead to temperature alarms if not corrected immediately.

The fan status is continuously monitored by the Cisco MDS SAN-OS software. In case of a fan failure, the following action is taken:

- System messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).

Use the **show environment fan** command to display the fan module status (see Example 10-6).

Example 10-6 Displays Chassis Fan Information

```
switch# show environment fan
----------------------------------------------------
Fan             Model                Hw         Status
----------------------------------------------------
Chassis         DS-9SLOT-FAN         1.2        ok
PS-1            --                   --         ok
PS-2            --                   --         absent
```

The possible Status field values for a fan module on the Cisco MDS 9500 Series switches are as follows:

- If the fan module is operating properly, the status is ok.
- If the fan is physically absent, the status is absent.
- If the fan is physically present but not working properly, the status is failure.

On the Cisco MDS 9513 Director, the front fan module has 15 fans. If the front fan module (DS-13SLT-FAN-F) State field contains “failure” in the **show environment fan** command output, it also displays the numbers of the failing fans (see Example 10-7).

Example 10-7 Displays Cisco MDS 9513 Front Fan Module Failure

```
switch# show environment fan
----------------------------------------------------
Fan             Model                Hw         Status
----------------------------------------------------
Chassis         DS-13SLT-FAN-F       0.3        failure 3 5 6 13
Chassis         DS-13SLT-FAN-R       0.3        ok
PS-1            --                   --         ok
PS-2            --                   --         ok
```

Figure 10-1 shows the numbering of the fans in the front fan module on the Cisco MDS 9513 Director.

Figure 10-1 Cisco MDS 9513 Front Fan Module Numbering

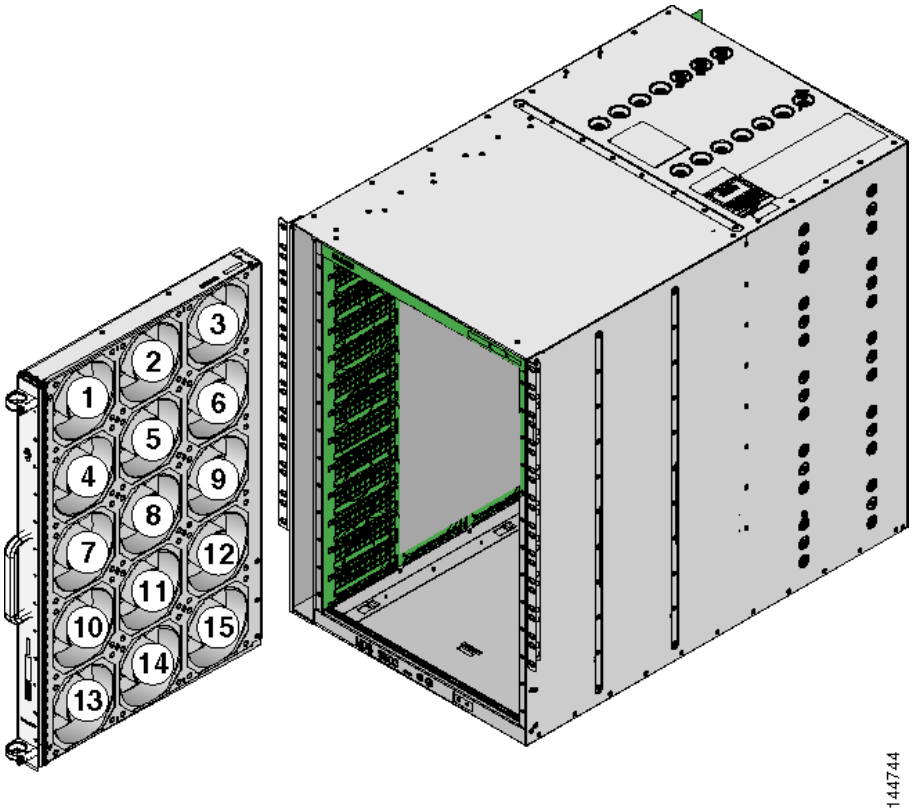

The rear fan module (DS-13SLT-FAN-R) on the Cisco MDS 9513 Director has only two fans. If a fan in the rear fan module fails, the State field in the **show environment fan** command output only displays "failure" and not the failing fan number (see Example 10-8).

Example 10-8 Displays Cisco MDS 9513 Rear Fan Module Failure

```
switch# show environment fan
----------------------------------------------------
Fan             Model                Hw         Status
----------------------------------------------------
Chassis         DS-13SLT-FAN-F       0.3        ok
Chassis         DS-13SLT-FAN-R       0.3        failure
PS-1            --                   --         ok
PS-2            --                   --         ok
```

About Clock Modules

All switches in the Cisco MDS 9000 Family have two clock modules—Module A (primary) and Module B (redundant). The clock modules are designed, tested, and qualified for mission-critical availability with a mean time between failures (MTBF) of 3,660,316 hours. This translates to a potential failure every 365 years. Additionally, Cisco MDS 9000 Family switches are designed to automatically switch to the redundant clock module should the active clock module fail.

Tip We recommend that the failed clock module be replaced during a maintenance window.

Use the **show environment clock** command to display the status for both clock modules (see Example 10-9).

Example 10-9 Displays Chassis Clock Information

```
switch# show environment clock
-------------------------------------------------------
Clock           Model                Hw         Status
-------------------------------------------------------
A               DS-C9500-CL          0.0        ok/active
B               DS-C9500-CL          0.0        ok/standby
```

Displaying Environment Information

Use the **show environment** command to display all environment-related switch information.

Example 10-10 Displays All Environment Information

```
switch# show environment
Clock:
-------------------------------------------------------
Clock           Model                Hw         Status
-------------------------------------------------------
A               Clock Module         1.0        ok/active
B               Clock Module         1.0        ok/standby

Fan:
-----------------------------------------------------
FAN             Model                Hw         Status
-----------------------------------------------------
Chassis         DS-2SLOT-FAN         0.0        ok
PS-1            --                   --         ok
PS-2            --                   --         absent
Temperature:
------------------------------------------------------------------
Module   Sensor   MajorThresh   MinorThres   CurTemp     Status
                  (Celsius)     (Celsius)    (Celsius)
------------------------------------------------------------------
1        1        75            60           32          ok
1        2        65            50           32          ok
1        3        -127          -127         43          ok
1        4        -127          -127         39          ok
Power Supply:
-----------------------------------------------------
PS  Model                Power     Power       Status
                         (Watts)   (Amp @42V)
-----------------------------------------------------
1   PWR-950-AC           919.38    21.89       ok
2                        --        --          absent
Mod Model                Power     Power       Power     Power       Status
                         Requested Requested   Allocated Allocated
                         (Watts)   (Amp @42V)  (Watts)   (Amp @42V)
--- -------------------  -------   ----------  --------- ----------  ----------
1    DS-X9216-K9-SUP     220.08    5.24        220.08    5.24        powered-up
```

```
Power Usage Summary:
--------------------
Power Supply redundancy mode:                 redundant
Total Power Capacity                          919.38   W
Power reserved for Supervisor(s)[-]           220.08   W
Power reserved for Fan Module(s)[-]           0.00     W
Power currently used by Modules[-]            0.00     W
                                              -------
Total Power Available                         699.30   W
                                              -------
```

Default Settings

Table 10-4 lists the default hardware settings.

Table 10-4 Default Hardware Parameters

Parameters	Default
Power supply mode	Redundant mode.

CHAPTER 11

Managing Modules

This chapter describes how to manage switching and services modules (also known as line cards) and provides information on monitoring module states.

This chapter includes the following sections:

About Modules

Table 11-1 describes the supervisor module options for switches in the Cisco MDS 9000 Family.

Table 11-1 Supervisor Module Options

Product	Number of Supervisor Modules	Supervisor Module Slot Number	Switching and Services Module Features
Cisco MDS 9513	Two modules	7 and 8	13-slot chassis allows any switching or services module in the other eleven slots.
Cisco MDS 9509	Two modules	5 and 6	9-slot chassis allows any switching or services module in the other seven slots.
Cisco MDS 9506	Two modules	5 and 6	6-slot chassis allows any switching or services module in the other four slots.
Cisco MDS 9216	One module	1	2-slot chassis allows one optional switching or services module in the other slot.
Cisco MDS 9216A	One module	1	2-slot chassis allows one optional switching or services module in the other slot.
Cisco MDS 9216i	One module	1	2-slot chassis allows one optional switching or services module in the other slot.

Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

- Cisco MDS 9513 Directors have two supervisor modules—one in slot 7 (sup-1) and one in slot 8 (sup-2). See Table 11-2. When the switch powers up and both supervisor modules come up together, the active module is the one that comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.
- Cisco MDS 9506 and Cisco MDS 9509 switches have two supervisor modules—one in slot 5 (sup-1) and one in slot 6 (sup-2). See Table 11-2. When the switch powers up and both supervisor modules come up together, the active module is the one that comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.
- Cisco MDS 9216i switches have one supervisor module that includes an integrated switching module with 14 Fibre Channel ports and two Gigabit Ethernet ports.
- Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.

Table 11-2 Supervisor Module Terms and Usage in Console Displays

Module Terms	Fixed or Relative	Usage
`module-7 and module-8`	Fixed usage for MDS 9513	module-7 always refers to the supervisor module in slot 7 and module-8 always refers to the supervisor module in slot 8.
`module-5 and module-6`	Fixed usage for MDS 9509 and MDS 9506	module-5 always refers to the supervisor module in slot 5 and module-6 always refers to the supervisor module in slot 6.
`module-1`	Fixed usage for MDS 9200 series	module-1 always refers to the supervisor module in slot 1.
`sup-1 and sup-2`	Fixed usage	On the MDS 9506 and MDS 9509 switches, sup-1 always refers to the supervisor module in slot 5 and sup-2 always refers to the supervisor module in slot 6. On the MDS 9513 Directors, sup-1 always refers to the supervisor module in slot 7 and sup-2 always refers to the supervisor module in slot 8.
`sup-active and sup-standby`	Relative usage	sup-active refers to the active supervisor module—relative to the slot that contains the active supervisor module. sup-standby refers to the standby supervisor module—relative to the slot that contains the standby supervisor module.
`sup-local and sup-remote`	Relative usage	If you are logged into the active supervisor, sup-local refers to the active supervisor module and sup-remote refers to the standby supervisor module. If you are logged into the standby supervisor, sup-local refers to the standby supervisor module (the one you are logged into.) There is no sup-remote available from the standby supervisor module (you cannot access a file system on the active sup).

Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. These modules obtain their image from the supervisor module.

Services Modules

Cisco MDS 9000 Family switches support any services module in any non-supervisor slot.

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

Verifying the Status of a Module

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time issue the **show module** command(see the "Fibre Channel Interfaces" section on page 12-1). The interfaces in each module are ready to be configured when the `ok` status is displayed in the **show module** command output. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
2    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    0      Caching Services Module                             ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
8    0      Caching Services Module          DS-X9560-SMAP      ok
9    32     1/2 Gbps FC Module               DS-X9032           ok

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
2    1.3(0.106a)  0.206   20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a)  0.602   --
6    1.3(0.106a)) 0.602   -- <---------------- New running version in module 6
8    1.3(0.106a)  0.702   --
9    1.3(0.106a)  0.3     22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de  JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd  JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe  JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6  JAB06280ae9

* this terminal session
```

The Status column in the output should display an *ok* status for switching modules and an active or standby (or HA-standby) status for supervisor modules. If the status is either ok or active, you can continue with your configuration..

Note A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled (see the "HA Switchover Characteristics" section on page 9-2). If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

The states through which a switching module progresses is discussed in the "Checking the State of a Module" section on page 11-4.

Checking the State of a Module

If your chassis has more than one switching module (also known as line card), you can check the progress by issuing the **show module** command several times and viewing the `Status` column each time.

The switching module goes through a testing and an initializing stage before displaying an `ok` status. Table 11-3 describes the possible states in which a module can exist.

Table 11-3 Module States

Module Status Output	Description
`powered up`	The hardware has electrical power. When the hardware is powered up, the software begins booting.
`testing`	The switching module has established connection with the supervisor module and the switching module is performing bootup diagnostics.
`initializing`	The diagnostics have completed successfully and the configuration is being downloaded.
`failure`	The switch detects a switching module failure upon initialization and automatically attempts to power-cycle the module three times. After the third attempt it continues to display a failed state.
`ok`	The switch is ready to be configured.
`power-denied`	The switch detects insufficient power for a switching module to power up.
`active`	This module is the active supervisor module and the switch is ready to be configured.
`HA-standby`	The HA switchover mechanism is enabled on the standby supervisor module (see the "HA Switchover Characteristics" section on page 9-2).
`standby`	The warm switchover mechanism is enabled on the standby supervisor module (see the "HA Switchover Characteristics" section on page 9-2).

Connecting to a Module

At any time, you can connect to any module using the **attach module** command. Once you are at the module prompt, you can obtain further details about the module using module-specific commands in EXEC mode.

To attach to a module, follow these steps:

	Command	Purpose
Step 1	`switch# attach module 6` `switch(standby)#`	Provides direct access to the specified module (in this example, the standby supervisor module is in slot 6).

Step 2	`switch(standby)# dir bootflash:`
	
	`root 14502912 Jan 13 12:23:52 1980 kickstart_image1`
	`admin 14424576 Jan 14 06:47:29 1980 kickstart_image2`
	`admin 14469632 Jan 14 01:29:16 1980 kickstart_image3`
	`root 14490112 Jan 08 07:25:50 1980 kickstart_image4`
	`root 12288 Jan 16 15:49:24 1980 lost+found/`
	`admin 14466048 Jan 14 02:40:16 1980 kickstart_image5`
	`admin 24206675 Jan 14 02:57:03 1980 m9500-sf1ek.bin`
	`root 19084510 Jan 13 12:23:28 1980 system_image1`
	`admin 19066505 Jan 14 06:45:16 1980 system_image2`
	`admin 18960567 Jan 14 01:25:21 1980 system_image5`
	
	`Usage for bootflash: filesystem`
	`158516224 bytes total used`
	`102400 bytes free`
	`167255040 bytes available`
		Provides the available space information for the standby supervisor module.
		Note Type **exit** to exit the module-specific prompt.
		Tip If you are not accessing the switch from a console terminal, this is the only way to access the standby supervisor module.

You can also use the **attach module** command as follows:

- To display the standby supervisor module information, although you cannot configure the standby supervisor module using this command.
- To display the switching module portion of the Cisco MDS 9200 Series supervisor module which resides in slot 1.

Reloading Modules

You can reload the entire switch, reset specific modules in the switch, or reload the image on specific modules in the switch.

This section includes the following totopics:

Reloading a Switch

To reload the switch, issue the **reload** command without any options. When you issue this command, you reboot the switch (see Chapter 7, “Software Images”).

Note If you need to issue the **reload** command, be sure to save the running configuration using the **copy running-config startup-config** command.

Power Cycling Modules

To power cycle any module, follow these steps:

Step 1 Identify the module that needs to be reset.

Step 2 Issue the **reload module** command to reset the identified module. This command merely power cycles the selected module.

```
switch# reload module number
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch# reload module 2
```

Caution Reloading a module disrupts traffic through the module.

Reloading Switching Modules

Switching modules automatically download their images from the supervisor module and do not need a forced download. This procedure is provided for reference should a need arise.

To replace the image on a switching module, follow these steps:

Step 1 Identify the switching module that requires the new image.

Step 2 Issue the **reload module** *number* **force-dnld** command to update the image on the switching module.

```
switch# reload module number force-dnld
```

Where *number* indicates the slot in which the identified module resides. In this example, the identified module resides in slot 9.

```
switch# reload module 9 force-dnld...
Jan  1 00:00:46 switch %LC-2-MSG:SLOT9 LOG_LC-2-IMG_DNLD_COMPLETE: COMPLETED
downloading of linecard image. Download successful...
```

Caution Reloading a module disrupts traffic through the module.

Preserving Module Configuration

Issue the **copy running-config startup-config** command from EXEC mode to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

Table 11-4 displays various scenarios when module configurations are preserved or lost.

Table 11-4 Switching Module Configuration Status

Scenario	Consequence
A particular switching module is removed and the **copy running-config startup-config** command is issued again.	The configured module information is lost.
A particular switching module is removed and the same switching module is replaced before the **copy running-config startup-config** command is issued again.	The configured module information is preserved.
A particular switching module is removed and replaced with the same type switching module, and a **reload module** *number* command is issued.	The configured module information is preserved.
A particular switching module is reloaded when a **reload module** *number* command is issued.	The configured module information is preserved.
A particular switching module is removed and replaced with a different type of switching module. For example, a 16-port switching module is replaced with a 32-port switching module. Sample scenario: 1. The switch currently has a 16-port switching module and the startup and running configuration files are the same. 2. You replace the 16-port switching module in the switch with a 32-port switching module. 3. Next, you remove the 32-port switching module and replace it with the same 16-port switching module referred to in Step 1. 4. You **reload** the switch.	The configured module information is lost from the running configuration. The default configuration is applied. The configured module information remains in startup configuration until a **copy running-config startup-config** command is issued again. Sample response: 1. The switch uses the 16-port switching module and the present configuration is saved in nonvolatile storage. 2. The factory default configuration is applied. 3. The factory default configuration is applied. 4. The configuration saved in nonvolatile storage referred to in Step 1 is applied.

Purging Module Configuration

Issue the **purge module** *slot* **running-config** command from EXEC mode to delete the configuration in a specific module. Once this command is issued, the running configuration is cleared for the specified slot. This command does not work on supervisor modules or on any slot that currently has a module. This command only works on an empty slot (where the specified module once resided).

The **purge module** command clears the configuration for any module that previously existed in a slot and has since been removed. While the module was in that slot, some parts of the configuration may have been stored in the running configuration and cannot be reused (for example, IP addresses), unless it is cleared from the running configuration.

For example, suppose you create an IP storage configuration with an IPS module in slot 3 in Switch A. This module uses IP address 10.1.5.500. You decide to remove this IPS module and move it to Switch B, and you no longer need the IP address10.1.5.500. If you try to configure this unused IP address, you

will receive an error message that prevents you from proceeding with the configuration. In this case, you need to issue the **purge module 3 running-config** command to clear the old configuration in Switch A before proceeding with using this IP address.

Powering Off Switching Modules

By default, all switching modules are in the power up state.

To power off a module, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# poweroff module 1` `switch(config)#`	Powers off the specified module (switching module 1) in the switch.
	`switch(config)# no poweroff module 1` `switch(config)#`	Powers up the specified module (switching module 1) in the switch.

Identifying Module LEDs

Table 11-5 describes the LEDs for the Cisco MDS 9200 Series integrated supervisor modules.

Table 11-5 LEDs for the Cisco MDS 9200 Series Supervisor Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage. The system will be shut down after two minutes if this condition is not cleared.
Speed[1]	On	2-Gbps mode and beacon mode disabled.
	Off	1-Gbps mode and beacon mode disabled.
	Flashing	Beacon mode enabled. See the "Identifying the Beacon LEDs" section on page 12-17.

Table 11-5 LEDs for the Cisco MDS 9200 Series Supervisor Modules (continued)

LED	Status	Description
Link	Solid green	Link is up.
	Solid yellow	Link is disabled by software.
	Flashing yellow	A fault condition exists.
	Off	No link.

1. The speed LED is only present on the 2-Gbps Fibre Channel switching modules.

Table 11-6 describes the LEDs for the Cisco MDS 9200 Series interface module.

Table 11-6 LEDs on the Cisco MDS 9200 Series Interface Module

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage.
System	Green	All chassis environmental monitors are reporting OK.
	Orange	The power supply failed or the power supply fan failed. or Incompatible power supplies are installed. or The redundant clock failed.
	Red	The temperature of the supervisor module exceeded the major threshold.
MGMT 10/100 Ethernet Link LED	Green	Link is up.
	Off	No link.
MGMT 10/100 Ethernet Activity LED	Green	Traffic is flowing through port.
	Off	No link or no traffic.

Table 11-7 describes the LEDs for the 16-port and 32-port switching modules, and the 4-port, 12-port, 24-port, and 48-port Generation 2 switching modules.

Table 11-7 LEDs for the Cisco MDS 9000 Family Fibre Channel Switching Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation.
	Red	The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage.
Speed	On	2-Gbps mode.
	Off	1-Gbps mode.
Link	Solid green	Link is up.
	Steady flashing green	Link is up (beacon used to identify port).
	Intermittent flashing green	Link is up (traffic on port).
	Solid yellow	Link is disabled by software.
	Flashing yellow	A fault condition exists.
	Off	No link.

The LEDs on the supervisor module indicate the status of the supervisor module, power supplies, and the fan module. Table 11-8 provides more information about these LEDs.

Table 11-8 LEDs for the Cisco MDS 9500 Series Supervisor Modules

LED	Status	Description
Status	Green	All diagnostics pass. The module is operational (normal initialization sequence).
	Orange	The module is booting or running diagnostics (normal initialization sequence). or An over temperature condition has occurred (a minor threshold has been exceeded during environmental monitoring).
	Red	The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or An over temperature condition occurred (a major threshold was exceeded during environmental monitoring).
System[1]	Green	All chassis environmental monitors are reporting OK.
	Orange	The power supply has failed or the power supply fan has failed. or Incompatible power supplies are installed. or The redundant clock has failed.
	Red	The temperature of the supervisor module major threshold has been exceeded.
Active	Green	The supervisor module is operational and active.
	Orange	The supervisor module is in standby mode.
Pwr Mgmt[1]	Green	Sufficient power is available for all modules.
	Orange	Sufficient power is not available for all modules.
MGMT 10/100 Ethernet Link LED	Green	Link is up.
	Off	No link.
MGMT 10/100 Ethernet Activity LED	Green	Traffic is flowing through port.
	Off	No link or no traffic.
CompactFlash	Green	The external CompactFlash card is being accessed.
	Off	No activity.

1. The System and Pwr Mgmt LEDs on a redundant supervisor module are synchronized to the active supervisor module.

EPLD Configuration

Switches and directors in the Cisco MDS 9000 Family contain several electrical programmable logical devices (EPLDs) that provide hardware functionalities in all modules. EPLD image upgrades are periodically provided to include enhanced hardware functionality or to resolve known issues.

Tip Refer to the *Cisco MDS SAN-OS Release Notes* to verify if the EPLD has changed for the Cisco SAN-OS image version being used.

EPLDs can be upgraded or downgraded using CLI commands. When EPLDs are being upgraded or downgraded, the following guidelines and observations apply:

- You can individually update each module that is online. The EPLD update is only disruptive to the module being upgraded.
- If you interrupt an upgrade, the module must be upgraded again.
- The upgrade or downgrade can only be executed from the active supervisor module. While the active supervisor module cannot be updated, you can update the other modules individually.
- In Cisco MDS 9100 Series fabric switches, be sure to specify one (1) as the module number.
- Cisco MDS 9200 Series switches do not support EPLD upgrades.
- The upgrade and downgrade processes disrupt traffic.

Caution Do not insert or remove any modules while an EPLD upgrade or downgrade is in progress.

Upgrading EPLD Images

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Issue the **show version** command to verify the Cisco MDS SAN-OS release running on the MDS switch.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2006, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software may be covered under the GNU Public
License or the GNU Lesser General Public License. A copy of
each such license is available at
http://www.gnu.org/licenses/gpl.html and
http://www.gnu.org/licenses/lgpl.html

Software
  BIOS:      version 1.0.8
  loader:    version unavailable [last: 1.0(0.267c)]
  kickstart: version 2.1(2) [build 2.1(2.47)] [gdb]
  system:    version 2.1(2) [build 2.1(2.47)] [gdb]

...
```

Step 3 If necessary, upgrade the Cisco MDS SAN-OS software running on your switch (see the "Software Upgrade Methods" section on page 7-6).

Step 4 Issue the **dir bootflash:** or **dir slot0:** command to verify that the EPLD software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(2), you must have m9000-epld-2.1.2.img in bootflash: or slot0: on the active supervisor module.

```
switch# dir bootflash:
  12288 Jan 01 00:01:07 1980 lost+found/
2337571 May 31 13:43:02 2005 m9000-epld-2.1.2.img
...
```

You can find the EPLD images at the following URL:

http://www.cisco.com/cgi-bin/tablebuild.pl/mds-epld

Step 5 If you need to obtain the appropriate EPLD software image file, follow these steps:

a. Download the EPLD software image file from Cisco.com to your FTP server.

b. Verify that you have enough free space available on the active and standby supervisor memory devices that you plan to use, either bootflash: or slot0:. The download site on Cisco.com shows the size of the EPLD image file in bytes.

The following example shows how to display the available memory for the bootflash: devices on the active and standby supervisors.

```
switch# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
2    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total
```

```
switch(standby)# exit
switch#
```

The following example shows how to display the available memory for the slot0: devices on the active and standby supervisors.

```
switch# dir slot0:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for slot:
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
2    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# dir slot0:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for slot0:
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch(standby)# exit
switch#
```

c. If there is not enough space, delete unneeded files.

```
switch# delete bootflash:m9500-sf1ek9-kickstart-mz.2.1.1.bin
switch# attach module 6
switch(standby)#
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch(standby)# delete bootflash:m9500-sf1ek9-kickstart-mz.2.1.1.bin
switch(standby)# exit
switch#
```

d. Copy the EPLD image file from the FTP server to the bootflash: or slot0: device in the active supervisor module. The following example shows how to copy to bootflash:

```
switch# copy ftp://10.1.7.2/m9000-epld-2.1.2.img bootflash:m9000-epld-2.1.2.img
```

Note The system will automatically synchronize the ELPD image to the standby supervisor if automatic copying is enabled.

```
switch# config t
switch(config)# boot auto-copy
```

Step 6 Use the **install module** *number* **epld** *url* command on the active supervisor module to upgrade EPLD images for a module.

```
switch# install module 2 epld bootflash:m9000-epld-2.1.2.img

EPLD                                  Curr Ver    New Ver
------------------------------------------------------------
XBUS IO                               0x07        0x07
UD Flow Control                       0x05        0x05
PCI ASIC I/F                          0x05        0x05
PCI Bridge                            0x05        0x07
WARNING: Upgrade process could take upto 15 minutes.

Module 2 will be powered down now!!
Do you want to continue (y/n) ? y
\ <----------------------------------------------------------progress twirl
Module 2 EPLD upgrade is successful
```

If you forcefully upgrade a module that is not online, all EPLDs are forcefully upgraded. If the module is not present in the switch, an error is returned. If the module is present, the command process continues. To upgrade a module that is not online but is present in the chassis, use the same command. The switch software prompts you to continue after reporting the module state. When you confirm your intention to continue, the upgrade continues.

```
switch# install module 2 epld bootflash:m9000-epld-2.1.2.img
\ <----------------------------------------------------------progress twirl
Module 2 EPLD upgrade is successful
```

Note When you upgrade the EPLD module on Cisco MDS 9100 Series switches, you receive the following message:

```
Data traffic on the switch will stop now!!
Do you want to continue (y/n) ?
```

Note The same procedure used to upgrade the EPLD images on a module can be used to downgrade the EPLD images.

Displaying EPLD Versions

Use the **show version module** *number* **epld** command to view all current EPLD versions on a specified module (see Example 11-1).

Example 11-1 Displays Current EPLD Versions for a Specified Module

```
switch# show version module 2 epld
EPLD Device                          Version
---------------------------------------
Power Manager                        0x07
XBUS IO                              0x07
UD Flow Control                      0x05
PCI ASIC I/F                         0x05
PCI Bridge                           0x07
```

Use the **show version epld** *url* command to view the available EPLD versions (see Example 11-2).

Example 11-2 Displays Available EPLD Versions

```
switch# show version epld bootflash:m9000-epld-2.1.1a.img
MDS series EPLD image, built on Wed May  4 09:52:37 2005

Module Type                          EPLD Device          Version
------------------------------------------------------------------
MDS 9500 Supervisor 1                XBUS 1 IO               0x09
                                     XBUS 2 IO               0x0c
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x04

1/2 Gbps FC Module (16 Port)         XBUS IO                 0x07
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05

1/2 Gbps FC Module (32 Port)         XBUS IO                 0x07
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05

Advanced Services Module             XBUS IO                 0x07
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05
                                     PCI Bridge              0x07

IP Storage Services Module (8 Port)  Power Manager           0x07
                                     XBUS IO                 0x03
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05
                                     Service Module I/F      0x0a
                                     IPS DB I/F              0x1a

IP Storage Services Module (4 Port)  Power Manager           0x07
                                     XBUS IO                 0x03
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05
                                     Service Module I/F      0x1a

Caching Services Module              Power Manager           0x08
                                     XBUS IO                 0x03
                                     UD Flow Control         0x05
                                     PCI ASIC I/F            0x05
                                     Service Module I/F      0x72
```

```
                                    Memory Decoder 0        0x02
                                    Memory Decoder 1        0x02

MDS 9100 Series Fabric Switch       XBUS IO                 0x03
                                    PCI ASIC I/F            0x40000003

2x1GE IPS, 14x1/2Gbps FC Module     Power Manager           0x07
                                    XBUS IO                 0x05
                                    UD Flow Control         0x05
                                    PCI ASIC I/F            0x07
                                    IPS DB I/F              0x1a
```

SSM Feature Support

Table 11-9 lists the features supported on the Cisco MDS SAN-OS Release 2.x for the SSM.

Table 11-9 Cisco MDS SAN-OS Release 2.x Feature Support for SSMs.

Cisco MDS SAN-OS Release			
2.0(1b)	**2.0(2b), 2.0(3), 2.0(4), and 2.0(4a)**	**2.1(1a)**	**2.1(2)and later**
None	Fibre Channel switching Intelligent Storage Services VSFN	Fibre Channel switching Intelligent Storage Services VSFN	Fibre Channel switching Intelligent Storage Services Nondisruptive upgrade for Fibre Channel switching traffic[1]

1. Requires EPLD version 2.1(2). See "EPLD Configuration" section on page 11-13.

Installing the SSI Boot Image on an SSM

This section describes how to install the SSI boot image for the Cisco MDS 9000 Family 32-port Fibre Channel Storage Services Module (SSM). The SSM supports normal Fibre Channel switching and Intelligent Storage Services. To use Fibre Channel switching and Intelligent Storage Services, you must install an SSI boot image on the SSM.

Note A newly installed SSM initially operates in Fibre Channel switching mode by default.

To install the SSI boot image on an SSM, follow these steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Issue the **dir modflash://***slot***-1/** command to verify that the SSI boot image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco SAN-OS Release 2.1(2), you must have m9000-ek9-ssi-mz.2.1.2.bin in modflash: on the SSM. To determine the correct SSI boot image to use, refer to the *Cisco MDS SAN-OS Release Compatibility Matrix for Storage Service Interface Images*.

You can find the SSI images at the following URL:
http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-ssi-3des

Step 3 If the file is not present in bootflash: or the modflash:, follow these steps:

a. Issue the **dir modflash://***slot***-1/** command to ensure that there is enough free space for the SSI image file. If necessary, issue the **delete modflash://***slot***-1/***filename* command to remove files.

b. Download the appropriate SSI boot image file to your FTP server and copy it from an FTP server to modflash: on the SSM:

```
switch# copy ftp://10.1.7.2/m9000-ek9-ssi-mz.2.1.2.bin
modflash://4-1/m9000-ek9-ssi-mz.2.1.2.bin
```

Step 4 Use the **install ssi** command to install the SSI boot image on the SSM.

Note As of Cisco SAN-OS Release 3.0(2), if the SSI boot image is located on bootflash: the **install ssi** command copies the SSI boot image to the modflash: on the SSM.

```
switch# install ssi modflash://4-1/m9000-ek9-ssi-mz.2.1.2.bin
```

Step 5 Issue the **show module** command to verify the status of the SSM.

```
switch# show module
Mod Ports Module-Type                       Model              Status
--- ----- --------------------------------- ------------------ ------------
4   32    Storage Services Module           DS-X9032-SSM       ok
...
Mod      Application Image Description     Application Image Version
-------- -----------------------------     -------------------------
4        SSI linecard image                2.1(2)
...
```

Upgrading the SSI Boot Image on an SSM

As of Cisco SAN-OS Release 2.0(2b), you can specify the SSI boot image for a Storage Services Module (SSM) to configure Fibre Channel switching and Intelligent Storage Services (see Chapter 47, "Configuring SCSI Flow Services and Statistics", Chapter 48, "Configuring Fibre Channel Write Acceleration", Chapter 49, "Configuring SANTap", and Chapter 50, "Configuring NASB"). Once you set the SSI image boot variable, you do not need to reset it for upgrades or downgrades to any Cisco MDS SAN-OS release that supports the SSI image.

Note If your switch is running Cisco MDS SAN-OS Release 2.1(2) or later, a newly installed SSM initially operates in Fibre Channel switching mode by default.

Note If you downgrade to a Cisco MDS SAN-OS release that does not support the SSM, you must power down the module. The boot variables for the SSM are lost.

SSI Boot Image Upgrade Considerations for the SSM

When you upgrade, or downgrade, the SSI boot image on an SSM, you might disrupt traffic through the module. Table 11-10 describes how updating the SSI boot image affects SSM traffic.

Table 11-10 SSI Boot Image Updating Affects on SSM Traffic

Cisco MDS SAN-OS Release	Traffic Type	Disrupts Traffic?
2.0(2b) through 2.1(1a)	All	Yes
2.1(2) and later	Layer 2 Fiber Channel switching only	No[1]
	Both Layer 2 Fiber Channel switching and Layer 3 Intelligent Storage Services (such as FCWA, NASB, SANTap, ISAPI virtualization)	Yes
	Layer 3 Intelligent Storage Services (such as FCWA, NASB, SANTap, ISAPI virtualization) only	Yes

1. Requires EPLD version 2.1(2). See "EPLD Configuration" section on page 11-13.

As shown in Table 11-10, layer 3 Intelligent Storage Services traffic is disrupted when you update the SSI boot image. If you have configured layer 3 Intelligent Storage Services on your SSM, we recommend that you shutdown these services before upgrading the SSI boot image. You can use dual fabric configuration to minimize the impact of shutting down layer 3 services.

To upgrade or downgrade the SSI boot image Fibre Channel switching and Intelligent Storage Services, perform the following steps:

Step 1 Verify that the correct SSI boot image is present on your switch (see the "Verifying the SSI Boot Image" section on page 11-21).

Step 2 Update the SSI boot image using one of the following methods:

a. If your switch is running Cisco MDS SAN-OS Release 2.0(1a) through Release 2.1(1a), configure the SSI boot variable to upgrade or downgrade the SSI boot image on the module (see the "Configuring the SSI Image Boot Variable" section on page 11-24).

b. Use the **install ssi** command to upgrade or downgrade the SSI boot image on the module (see the "Using the install ssi Command" section on page 11-26).

Note The SSM must be running EPLD version 2.1(2) to use the **install ssi** command. You must install the SSM on a Cisco MDS 9500 Series switch to upgrade the EPLD. See the "EPLD Configuration" section on page 11-13.

Verifying the SSI Boot Image

To verify that you have the correct Cisco MDS SAN-OS release and SSI boot image file on your switch, perform the following steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Issue the **show version** command to ensure that your switch is running Cisco MDS SAN-OS Release 2.1(1a) or later system and kickstart images.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2006, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software may be covered under the GNU Public
License or the GNU Lesser General Public License. A copy of
each such license is available at
http://www.gnu.org/licenses/gpl.html and
http://www.gnu.org/licenses/lgpl.html

Software
  BIOS:      version 1.0.8
  loader:    version unavailable [last: 1.0(0.267c)]
  kickstart: version 2.1(2) [build 2.1(2.47)] [gdb]
  system:    version 2.1(2) [build 2.1(2.47)] [gdb]

...
```

Step 3 If necessary, upgrade the Cisco MDS SAN-OS software running on your switch (see the "Software Upgrade Methods" section on page 7-6).

Step 4 Issue the **dir bootflash:** or **dir slot0:** command to verify that the SSI software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(2), you must have m9000-ek9-ssi-mz.2.1.2.bin in bootflash: or slot0: on the active supervisor module. Refer to the *Cisco MDS SAN-OS Release Compatibility Matrix for Storage Service Interface Images*.

Note As of Cisco MDS SAN-OS Release 2.1(2), we recommend that you use modflash: on the SSM. You can check for the presence of the SSI software image using the **dir modflash://***slot***-1/** command.

```
switch# dir bootflash:
  12288 Jan 01 00:01:07 1980 lost+found/
3821032 May 10 13:43:02 2005 m9000-ek9-ssi-mz.2.1.2.bin
...
```

You can find the SSI images at the following URL:

http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-ssi-3des

Step 5 If you need to obtain the appropriate SSI software image file, perform the following steps:

a. Download the SSI software image file from Cisco.com to your FTP server.

b. Verify that you have enough free space available on the active and standby supervisor memory devices which you plan to use, either bootflash: or slot0:. The download site on Cisco.com shows the size of the boot image file in bytes.

Note As of Cisco MDS SAN-OS Release 2.1(2), we recommend that you use modflash: on the SSM. You can check the available space using the **dir modflash://***slot***-1/** command.

The following example shows how to display the available memory for the bootflash: devices on the active and standby supervisors.

```
switch# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# dir bootflash:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for bootflash://sup-local
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch(standby)# exit
switch#
```

The following example shows how to display the available memory for the slot0: devices on the active and standby supervisors.

```
switch# dir slot0:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for slot:
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# dir slot0:
   12288 Jan 01 00:01:06 1980 lost+found/
14765056 Mar 21 15:35:06 2005 m9500-sf1ek9-kickstart-mz.2.1.1.bin
15944704 Apr 06 16:46:04 2005 m9500-sf1ek9-kickstart-mz.2.1.1a.bin
48063243 Mar 21 15:34:46 2005 m9500-sf1ek9-mz.2.1.1.bin
48036239 Apr 06 16:45:41 2005 m9500-sf1ek9-mz.2.1.1a.bin

Usage for slot0:
141066240 bytes used
 43493376 bytes free
184559616 bytes total

switch(standby)# exit
switch#
```

c. Delete the unneeded files, if there is not enough space.

```
switch# delete bootflash:m9500-sf1ek9-kickstart-mz.2.1.1.bin
```

The **show module** command output shows that the standby supervisor is in slot 6. Use the **attach** command to access the supervisor module.

```
switch# attach module 6
...
switch(standby)# delete bootflash:m9500-sf1ek9-kickstart-mz.2.1.1.bin
switch(standby)# exit
switch#
```

d. Copy the boot image file from the FTP server to the bootflash: or slot0: device in the active supervisor module. The following example shows how to copy to bootflash:

Note As of Cisco MDS SAN-OS Release 2.1(2), we recommend that you copy the image to modflash: on the SSM.

```
switch# copy ftp://10.1.7.2/m9000-ek9-ssi-mz.2.1.2.bin
bootflash:m9000-ek9-ssi-mz.2.1.2.bin
```

Note The system will automatically synchronize the SSI image to the standby supervisor if automatic copying is enabled.

```
switch# config t
switch(config)# boot auto-copy
```

Configuring the SSI Image Boot Variable

The following steps describe how to configure the SSI image boot variable for Fibre Channel switching and Intelligent Storage Services on the SSM on switches running Cisco MDS SAN-OS release 2.0(2b) through Release 2.1(1a) on the SSM.

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Verify that the SSM is physically installed in the switch. If the module is not physically installed, insert it into the desired slot. Issue a **show module** command to verify the status of the module.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

Note the slot number for later reference.

Step 3 Verify the Cisco MDS SAN-OS release running on the switch and the location and name of the SSI boot image on the switch following the procedure described in the "Verifying the SSI Boot Image" section on page 11-21.

Step 4 Configure the SSI image boot variable to specify the SSI image to use when the SSM reloads.

```
switch# config terminal
switch(config)# boot ssi bootflash:m9000-ek9-ssi-mz.2.1.1a.bin module 4
switch(config)# exit
switch#
```

Note You can only specify one image for the SSI variable per module.

Caution The SSI boot variable must reference the correct SSI boot image, otherwise the SSM fails to initialize. If you do not correctly set the SSI boot variable, the SSM remains in the power-down state after attempting to initialize three times.

Step 5 Issue the **show boot** command to display the current contents of the SSI boot variable.

```
switch# show boot
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
```

```
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
Module 4
ssi variable = bootflash:/m9000-ek9-ssi-mz.2.1.1a.bin
```

Step 6 Save the new boot variable configuration so the new boot image is used when the switch reboots.

```
switch# copy running-config startup-config
```

Note If you do not save this configuration, it is lost on a switch reboot. In addition the SSM stays in the power-down state if your switch is running Cisco MDS SAN-OS Release 2.1(1a) or earlier. You must perform this procedure again to recover the SSI image boot variable configuration.

Step 7 Reload the SSM to load the new boot image.

```
switch# reload module 4
reloading module 4 ...
```

The **reload** command power cycles the SSM.

Caution Reloading the SSM disrupts traffic through the module.

Step 8 Issue the **show module** command to verify the status of the SSM.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
4    2.1(2)       0.30    20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de
5    2.1(2)       4.0     --
6    2.1(2)       4.0     --

Mod      Application Image Description       Application Image Version
-------- ----------------------------        -------------------------
4        SSI linecard image                  2.1(1a)

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
4    00-05-30-00-9e-b2 to 00-05-30-00-9e-b6  JAB06480590
5    00-0e-38-c6-2c-6c to 00-0e-38-c6-2c-70  JAB082504MQ
6    00-0f-34-94-4d-34 to 00-0f-34-94-4d-38  JAB083407D3

* this terminal session
```

Using the install ssi Command

You can use the **install ssi** command to update the boot image on an SSM. If the SSM is performing Fibre Channel switching and no Intelligent Storage Services are provisioned on the module, this operation does not disrupt traffic through the module. If the SSM is configured for Intelligent Storage Services, a warning is displayed at the command prompt indicating that the operation will disrupt traffic and asking if you wish to continue.

Note The SSM must be running EPLD version 2.1(2) to use the **install ssi** command. You must install the SSM on a Cisco MDS 9500 Series switch to update the EPLD. See the "EPLD Configuration" section on page 11-13.

To upgrade or downgrade the SSM boot image for Intelligent Storage Services, follow these steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Verify that the SSM is physically installed in the switch. If the module is not physically installed, insert it into the desired slot. Issue a **show module** command to verify the status of the module.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok

5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
...
```

Note the slot number for later reference.

Step 3 Verify the Cisco MDS SAN-OS release running on the switch and the location and name of the SSI boot image on the switch following the procedure described in the "Verifying the SSI Boot Image" section on page 11-21.

Step 4 Install the SSI image on the SSM.

Note As of Cisco SAN-OS Release 3.0(2), if the SSI boot image is located on bootflash: the **install ssi** command copies the SSI boot image to the modflash: on the SSM.

```
switch# install ssi modflash://4-1/m9000-ek9-ssi-mz.2.1.2.bin module 4
```

Note If the SSM is configured for Layer 3 Fibre Channel switching or Intelligent Storage Services, a warning will be displayed at the command prompt indicating that the operation will disrupt traffic and you will be asked if you wish to continue.

Note As of Cisco MDS SAN-OS Release 2.1(2), we recommend that you reference the SSI boot image on modflash: on the SSM. Use the **install ssi modflash://***slot***-1/***filename* **module** *slot* command to install the SSI image.

Step 5 Issue the **show boot** command to display the current contents of the image boot variable for the SSM.

```
switch# show boot
```

```
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
Module 4
ssi variable = modflash://4-1/m9000-ek9-ssi-mz.2.1.2.bin
```

Step 6 Save the new boot variable configuration so the new boot image is used when the switch reboots.

```
switch# copy running-config startup-config
```

Note If you do not save this configuration, it is lost on a switch reboot. In addition, SSM stays in the power-down state if your switch is running Cisco MDS SAN-OS Release 2.1(1a) or earlier, or comes up in Fibre Channel switching mode if your switch is running Cisco MDS SAN-OS Release 2.1(2) or later. You must perform this procedure again to recover the SSI image boot variable configuration.

Step 7 Issue the **show module** command to verify the status of the SSM.

```
switch# show module
Mod  Ports  Module-Type                        Model              Status
---  -----  --------------------------------- ------------------ ------------
4    32     Storage Services Module           DS-X9032-SSM       ok
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
4    2.1(2)       0.30    20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de
5    2.1(2)       4.0     --
6    2.1(2)       4.0     --

Mod      Application Image Description       Application Image Version
-------- ----------------------------        -------------------------
4        SSI linecard image                  2.1(2)

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
4    00-05-30-00-9e-b2 to 00-05-30-00-9e-b6  JAB06480590
5    00-0e-38-c6-2c-6c to 00-0e-38-c6-2c-70  JAB082504MQ
6    00-0f-34-94-4d-34 to 00-0f-34-94-4d-38  JAB083407D3

* this terminal session
```

Managing SSMs and Supervisor Modules

This section describes the considerations for replacing SSMs and supervisor modules and for upgrading and downgrading Cisco MDS SAN-OS releases.

Considerations for Replacing SSMs and Supervisor Modules

If you replace an SSM or supervisor module, you should consider the following:

- If you replace an SSM with another SSM and the boot image is on bootflash:, respectively, you can leave the boot image installed on the active supervisor.
- If you replace an SSM with another SSM and the SSI boot image is on the modflash:, the SSM might not initialize. See the "Recovering an SSM After Replacing Corrupted CompactFlash Memory" section on page 11-28.
- If you replace an SSM with any other module, you can leave the boot image installed on the active supervisor or remove it. The active supervisor module detects the module type and boots the module appropriately.
- If you replace a supervisor module in a switch with active and standby supervisors, no action is required because the boot image is automatically synchronized to the new supervisor module.
- If you replace a supervisor module in a switch with no standby supervisor, you need to reimplement the configuration on the new supervisor.

Recovering an SSM After Replacing Corrupted CompactFlash Memory

In Cisco MDS SAN-OS Release 2.1(2) and later, you use the CompactFlash memory (modflash:) on the SSM to store the SSI image. If the modflash: on the SSM is replaced, the SSM might not initialize. To recover the SSM, follow these steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Display the values assigned to the SSI image boot variable for each module and note the values for later reference.

```
switch# show boot module
Module 2
ssi variable = modflash://2-1/m9000-ek9-ssi-mz.2.1.2.bin
Module 4
ssi variable = modflash://4-1/m9000-ek9-ssi-mz.2.1.2.bin
```

Step 3 Clear the values assigned to the SSI image boot variable.

```
switch# config t
switch(config)# no boot ssi
```

Step 4 Reload the SSM to initialize in Fibre Channel switching mode.

```
switch# reload module 4
reloading module 4 ...
```

Step 5 After the SSM initialize, follow the procedure described in the "Upgrading the SSI Boot Image on an SSM" section on page 11-19.

Step 6 Reassign the SSI boot variables cleared in Step 3,

```
switch# config t
```

```
switch(config)# boot ssi modflash://2-1/m9000-ek9-ssi-mz.2.1.2.bin module 2
```

Considerations for Upgrading and Downgrading Cisco MDS SAN-OS Releases

Consider the following when upgrading and downgrading the Cisco MDS SAN-OS software on a switch containing an SSM:

- Once you set the SSI image boot variable, you do not need to reset it for upgrades or downgrades to any Cisco MDS SAN-OS release that supports boot images. You can use the **install all** command or Fabric Manager GUI to upgrade SSMs once it has been installed. The CLI is required for the procedures described in the "Upgrading the SSI Boot Image on an SSM" section on page 11-19.
- If you downgrade to a Cisco MDS SAN-OS release that does not support the SSM, you must power down the module. The boot variables for the module are lost.
- The SSM cannot be configured for both the SSI and any other third-party software on the module such as VSFN.
- The following example shows successful **install all** command output including an SSI image.

Note The SSI boot variable setting is included in the **install all** output. Also, if the SSI boot image is located on bootflash: the **install all** command copies the SSI boot image to the modflash: on the SSMs.

```
Switch# install all system bootflash:isan-2-1-1a kickstart bootflash:boot-2-1-1a ssi
bootflash:ssi-2.1.1a

Copying image from bootflash:ssi-2.1.1a to modflash://2-1/ssi-2.1.1a.
[####################] 100% -- SUCCESS

Verifying image bootflash:/ssi-2.1.1a

[####################] 100% -- SUCCESS

Verifying image bootflash:/boot-2-1-1a

[####################] 100% -- SUCCESS

Verifying image bootflash:/isan-2-1-1a

[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-2-1-1a.

[####################] 100% -- SUCCESS

Extracting "ips4" version from image bootflash:/isan-2-1-1a.

[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-2-1-1a.

[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-2-1-1a.

[####################] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-2-1-1a.

[####################] 100% -- SUCCESS

Compatibility check is done:

Module bootable Impact Install-type Reason
------ -------- -------------- ------------ ------
2 yes non-disruptive rolling
3 yes disruptive rolling Hitless upgrade is not supported
4 yes disruptive rolling Hitless upgrade is not supported
5 yes non-disruptive reset

Images will be upgraded according to following table:

Module Image Running-Version New-Version Upg-Required
------ ---------- ------------------- -------------------- ------------
2 slc 2.0(3) 2.1(1a) yes
2 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
3 slc 2.0(3) 2.1(1a) yes
3 SSI 2.0(3) 2.1(1a) yes
3 bios v1.0.8(08/07/03) v1.1.0(10/24/03) yes
4 ips4 2.0(3) 2.1(1a) yes
4 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
5 system 2.0(3) 2.1(1a) yes
5 kickstart 2.0(3) 2.1(1a) yes
5 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
5 loader 1.2(2) 1.2(2) no

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Module 6:Force downloading.

-- SUCCESS

Syncing image bootflash:/SSI-2.1.1a to standby.

[####################] 100% -- SUCCESS

Syncing image bootflash:/boot-2-1-1a to standby.

[####################] 100% -- SUCCESS

Syncing image bootflash:/isan-2-1-1a to standby.

[####################] 100% -- SUCCESS

Setting boot variables.

[####################] 100% -- SUCCESS

Performing configuration copy.

[####################] 100% -- SUCCESS

Module 3:Upgrading Bios/loader/bootrom.

[####################] 100% -- SUCCESS

Module 6:Waiting for module online.
```

```
-- SUCCESS

"Switching over onto standby".

----------------------------
```

Default Settings

Table 11-11 lists the default settings for the supervisor module.

Table 11-11 Default Supervisor Module Settings

Parameters	Default
Administrative connection	Serial connection.
Global switch information	• No value for system name. • No value for system contact. • No value for location.
System clock	No value for system clock time.
In-band (VSAN 1) interface	IP address, subnet mask, and broadcast address assigned to the VSAN are set to 0.0.0.0.

Table 11-12 lists the default settings for the SSM.

Table 11-12 Default SSM Settings

Parameters	Default
Initial state when installed	• Power-down state on switches with Cisco MDS SAN-OS Release 2.1(1a) and earlier installed. • Fibre Channel switching mode on switches with Cisco MDS SAN-OS Release 2.1(2) and later installed and SSMs with EPLD version 2.0(2) and later installed.

PART 3

Switch Configuration

CHAPTER 12

Configuring Interfaces

A switch's main function is to relay frames from one data link to another. To do that, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces.

This chapter describes the basic interface configuration to get your switch up and running. It includes the following sections:

See Chapter 5, "Initial Configuration," and Chapter 43, "Configuring IP Services," for more information on configuring mgmt0 interfaces.

See Chapter 45, "Configuring IPv4 for Gigabit Ethernet Interfaces" and Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces"for more information on configuring Gigabit Ethernet interfaces.

Tip Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode (see the "Verifying the Module Status" section on page 5-16).

Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but not limited to) modes, frame encapsulation, states, SFPs, and speeds.

This section includes the following topics:

32-Port Switching Module Configuration Guidelines

The 32-port switching module guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain shutdown.
- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The 32-port switching module does not support FICON.

Note We recommend that you configure your E ports on a 16-port switching module. If you must configure an E port on a 32-port host optimized switching module, the other three ports in that 4-port group cannot be used.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, ST port, and B port (see Figure 12-1). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Figure 12-1 Cisco MDS 9000 Family Switch Port Modes

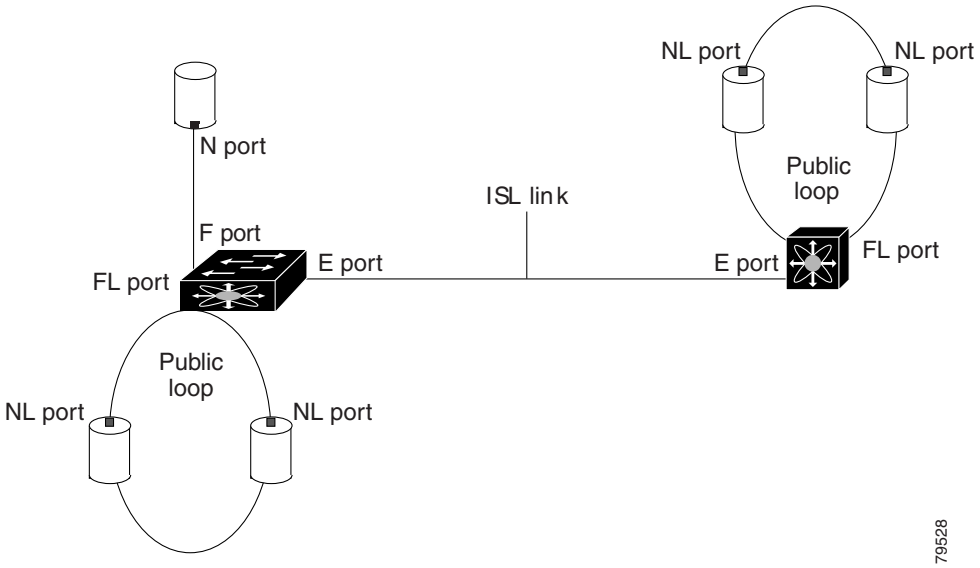

Note Interfaces are created in VSAN 1 by default. See Chapter 19, "Configuring and Managing VSANs."

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

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

Note When a module is removed and replaced with the same type of module, the configuration is retained. If a different type of module is inserted, then the original configuration is no longer retained.

A brief description of each interface mode follows.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see Chapter 16, "Configuring PortChannels").

Note We recommend that you configure E ports on 16-port modules. If you must configure an E port on a 32-port oversubscribed module, then you can only use the first port in a group of four ports (for example, ports 1 through 4, 5 through 8, and so forth). The other three ports cannot be used.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

Note FL port mode is not supported on 4-port 10-Gbps switching module interfaces.

NP Ports

An *NP port* is a port on a device that is in NPV mode and connected to the core switch via a F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

For more details about NP ports and NPV, see Chapter 13, "Configuring N Port Virtualization."

TL Port

In translative loop port (TL port) mode, an interface functions as a translative loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the "About TL Port ALPA Caches" section on page 12-30).

Tip We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

Note TL port mode is not supported on Generation 2 switching module interfaces.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see Chapter 15, "Configuring Trunking"). TE ports support class 2, class 3, and class F service.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes though a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they merely transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see Chapter 52, "Monitoring Network Traffic Using SPAN").

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the "Configuring SPAN" section on page 52-7).

Note ST port mode is not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

Figure 12-1 on page 12-3 depicts a typical SAN extension over an IP network.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see Chapter 44, "Configuring IP Storage").

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see Chapter 15, "Configuring Trunking").

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

Note Fibre Channel interfaces on Storage Services Modules (SSMs) cannot be configured in auto mode.

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. Figure 12-2 shows an example application using NPIV.

Figure 12-2 NPIV Example

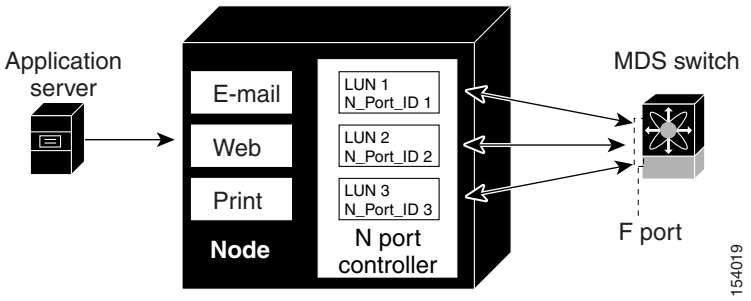

You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

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

About Interface States

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

Administrative States

The administrative state refers to the administrative configuration of the interface as described in Table 12-1.

Table 12-1 Administrative States

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

Operational States

The operational state indicates the current operational state of the interface as described in Table 12-2.

Table 12-2 Operational States

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

Reason Codes

Reason codes are dependent on the operational state of the interface as described in Table 12-3.

Table 12-3 Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.
Down	Down	Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	See Table 12-4.

Note Only some of the reason codes are listed in Table 12-4.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in Table 12-4.

Table 12-4 Reason Codes for Nonoperational States

Reason Code (long version)	Description	Applicable Modes
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The Cisco SAN-OS software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shut down or enable the interface.	
FC redirect failure	A port is isolated because a Fibre Channel redirect is unable to program routes.	
No port activation license available	A port is not active because it does not have a port license.	
SDM failure	A port is isolated because SDM is unable to program routes.	

Table 12-4 Reason Codes for Nonoperational States (continued)

Reason Code (long version)	Description	Applicable Modes
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to the other side of the link E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	

Configuring Fibre Channel Interfaces

To configure a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode. **Note** When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

To configure a range of interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1 - 4 , fc2/1 - 3` `switch(config-if)#`	Selects the range of Fibre Channel interfaces and enters interface configuration submode3. **Note** In this command, provide a space before and after the comma.

For the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter, you can configure a range of interfaces among internal ports or external ports, but you cannot mix both interface types within the same range. For example, "bay 1-10, bay 12" or "ext 0, ext 15-18" are valid ranges, but "bay 1-5, ext 15-17" is not.

Graceful Shutdown

Interfaces on a port are shutdown by default (unless you modified the initial configuration).

The Cisco SAN-OS software implicitly performs a graceful shutdown in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface.
- If a Cisco SAN-OS software application executes a port shutdown as part of its function.

A graceful shutdown ensures that no frames are lost when the interface is shutting down. When a shutdown is triggered either by you or the Cisco SAN-OS software, the switches connected to the shutdown link coordinate with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shutdown is not possible in the following situations:

- If you physically remove the port from the switch.
- If in-order-delivery (IOD) is enabled (see "In-Order Delivery" section on page 25-13).
- If the Min_LS_interval interval is higher than 10 seconds (see"Displaying Global FSPF Information" section on page 25-20).

Note This feature is only triggered if both switches at either end of this E port interface are MDS switches and are running Cisco SAN-OS Release 2.0(1b) or later.

Setting the Interface Administrative State

To gracefully shut down an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# shutdown`	Gracefully shuts down the interface and administratively disables traffic flow (default).

To enable traffic flow, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# no shutdown`	Enables traffic flow to administratively allow traffic when the **no** prefix is used (provided the operational state is up).

Configuring Interface Modes

To configure the interface mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport mode F` `switch(config-if)#`	Configures the administrative mode of the port. You can set the operational state to auto, E, F, FL, Fx, TL, or SD port mode. **Note** Fx ports refers to an F port or an FL port (host connection only), but not E ports.
	`switch(config-if)# switchport mode auto` `switch(config-if)#`	Configures the interface mode to auto-negotiate an E, F, FL, or TE port mode (not TL or SD port modes) of operation. **Note** TL ports and SD ports cannot be configured automatically. They must be administratively configured. **Note** You cannot configure Fibre Channel interfaces on SSMs in auto mode.

Configuring System Default Port Mode F

The **system default switchport mode F** command sets the administrative mode of all Fibre Channel ports to mode F, while avoiding traffic disruption caused by the formation of unwanted inter-switch links (ISLs). This command is part of the setup utility that runs during bootup after a **write erase** or **reload**. It can also be executed from the command line in configuration mode. This command changes the configuration of the following ports to administrative mode F:

- All ports that are down and that are not out-of-servicel.
- All F ports that are up, whose operational mode is F, and whose administrative mode is not F.

This command does not affect the configuration of the following ports:

- All user-configured ports, even if they are down.
- All non-F ports that are up; however, if non-F ports are down, this command changes the administrative mode of those ports.

Example 12-1 shows the command in the setup utility, and Example 12-2 shows the command from the command line.

Example 12-1 Setup Utility

```
Configure default switchport mode F (yes/no) [n]: y
```

Example 12-2 Command Line

```
switch(config)# system default switchport mode F
```

Note To ensure that ports that are part of ISLs do not get changed to port mode F, configure the ports in port mode E, rather than in Auto mode.

Note When the command is executed from the command line, switch operation remains graceful. No ports are flapped.

To sets the administrative mode of Fibre Channel ports to mode F in the CLI, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system default switchport mode F`	Sets the administrative mode of Fibre Channel ports to mode F (if applicable).
	`switch(config)# no system default switchport mode F`	Sets the administrative mode of Fibre Channel ports to the default (unless user configured).

Note For detailed information about the switch setup utility see Chapter 5, "Initial Configuration."

Configuring Port Speeds

By default, the portspeed for an interface is automatically calculated by the switch.

Caution Changing the port speed is a disruptive operation.

To configure the port speed of the interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1`	Selects the mgmt0 interface and enters interface configuration mode.
Step 3	`switch(config-if)# switchport speed 1000`	Configures the port speed of the interface to 1000 Mbps. The number indicates the speed in megabits per second (Mbps). You can set the speed to **1000** (for 1-Gbps interfaces), **2000** (for 2-Gbps interfaces), **4000** (for 4-Gbps interfaces), or **auto** (default).
	`switch(config-if)# no switchport speed`	Reverts the factory default (**auto**) administrative speed of the interface.

For *internal ports* on the Cisco Fabric Switch for HP c_Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter a port speed of 1 Gbps is not supported. Auto-negotiation is supported between 2 Gbps and 4 Gbps only. Also, if the BladeCenter is a "T" chassis, then port speeds are fixed at 2 Gbps and auto-negotiation is not enabled.

Autosensing

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

To avoid wasting unused bandwidth on 48-port and 24-port 4-Gbps Fibre Channel switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 Gbps. This feature shares the unused bandwidth within the port group provided that it does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports that are configured for autosensing.

Tip When migrating a host that supports up to 2-Gbps traffic (that is, not 4-Gbps with autosensing capabilities) to the 4-Gbps switching modules, use autosensing with a maximum bandwidth of 2-Gbps.

Enabling N Port Identifier Virtualization

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

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

To enable or disable NPIV on the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# npiv enable`	Enables NPIV for all VSANs on the switch.
Step 3	`switch(config)# no npiv enable`	Disables (default) NPIV on the switch.

About Interface Descriptions

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

Configuring the Interface Description

To configure a description for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport description cisco-HBA2`	Configures the description of the interface. The string can be up to 80 characters long.
	`switch(config-if)# no switchport description`	Clears the description of the interface.

About Frame Encapsulation

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

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (`Encapsulation is eisl`) in the **show interface** *SD_port_interface* command output (see the "Encapsulating Frames" section on page 52-10).

About Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Configuring Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure the receive data field size, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcrxbufsize 2000`	Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes.

Identifying the Beacon LEDs

Figure 12-3 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 12-3 Cisco MDS 9000 Family Switch Interface Modes

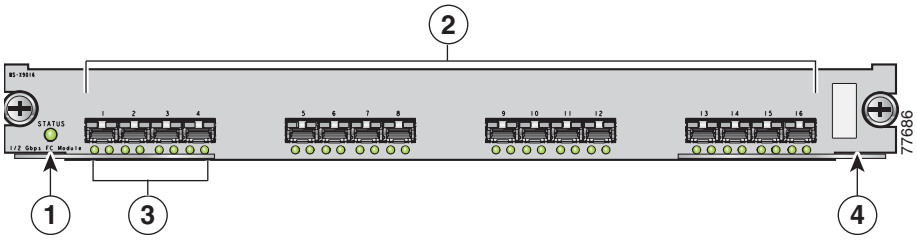

1	Status LED[1]	**3**	Link LEDs[1] and speed LEDs[2]
2	1/2-Gbps Fibre Channel port group[3]	**4**	Asset tag[4]

1. See the "Identifying Module LEDs" section on page 11-9.
2. See the "About Speed LEDs" section on page 12-17.
3. See the "32-Port Switching Module Configuration Guidelines" section on page 12-2.
4. Refer to the Cisco MDS 9000 Family hardware installation guide for your platform.

About Speed LEDs

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off or solid green—Beacon mode is disabled.
- Flashing green—The beacon mode is enabled. The LED flashes at one-second intervals.

About Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface.

Configuring the beacon mode has no effect on the operation of the interface.

Configuring Beacon Mode

To enable beacon mode for a specified interface or range of interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport beacon`	Enables the beacon mode for the interface.
	`switch(config-if)# no switchport beacon`	Disables the beacon mode for the interface.

Note The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

About Bit Error Thresholds

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

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- GBIC or SFP is specified to operate at 1 Gbps but is used at 2 Gbps
- GBIC or SFP is specified to operate at 2 Gbps but is used at 4 Gbps
- Short haul cable is used for long haul or long haul cable is used for short haul.
- Momentary sync loss
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached. You can issue **shutdown/no shutdown** command sequence to reenable the interface.

You can configure the switch to not disable an interface when the threshold is crossed. By default, the threshold disables the interface.

To disable the bit error threshold for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport ignore bit-errors`	Prevents the detection of bit error threshold events from disabling the interface.
	`switch(config-if)# no switchport ignore bit-errors`	Prevents the detection of bit error threshold events from enabling the interface.

Note Regardless of the setting of the **switchport ignore bit-errors** command, the switch generates a syslog message when bit error threshold events are detected.

Switch Port Attribute Default Values

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

To configure switch port attributes, follow these steps:

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

About SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fc***slot*/*port* **transceiver** command display both values for Cisco supported SFPs. Table 12-5 defines the acronyms used in the command output (see the "Displaying Interface Information" section on page 12-20).

Table 12-5 SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
short wave laser	swl
long wave laser	lwl
long wave laser cost reduced	lwcr
electrical	elec
Extended transmitters assigned to Cisco-supported SFPs	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530
CWDM-1550	c1550
CWDM-1570	c1570
CWDM-1590	c1590
CWDM-1610	c1610

Displaying Interface Information

The **show interface** command is invoked from the EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch. See Examples 12-3 to 12-10.

Example 12-3 Display All Interfaces

```
switch# show interface
fc1/1 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:0b:00:05:30:00:8d:de
    Admin port mode is F
    Port mode is F, FCID is 0x610000
    Port vsan is 2
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      134 frames input, 8468 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      154 frames output, 46072 bytes
        0 discards, 0 errors
      1 input OLS, 1 LRR, 0 NOS, 0 loop inits
      1 output OLS, 0 LRR, 1 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

```
. . .
fc1/9 is trunking
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:09:00:05:30:00:97:9e
    Peer port WWN is 20:0b:00:0b:5f:a3:cc:00
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 255
    Receive B2B Credit is 255
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100,3000)
    Trunk vsans (up)                       (1,100,3000)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 280 bits/sec, 35 bytes/sec, 0 frames/sec
    5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
      4609939 frames input, 8149405708 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      4638491 frames output, 7264731728 bytes
        0 discards, 0 errors
      3 input OLS, 9 LRR, 1 NOS, 0 loop inits
      9 output OLS, 7 LRR, 1 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
fc1/13 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:0d:00:05:30:00:97:9e
    Admin port mode is auto, trunk mode is on
    Port mode is F, FCID is 0x650100
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      8696 frames input, 3227212 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      16799 frames output, 6782444 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
sup-fc0 is up
    Hardware is Fibre Channel
    Speed is 1 Gbps
    139597 packets input, 13852970 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    139516 packets output, 16759004 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by issuing a command with the following example format:

```
interface fc1/1 - 5 , fc2/5 - 7
```

Note The spaces are required before and after the dash (**-**) and before and after the comma (**,**).

Example 12-4 Display Multiple, Specified Interfaces

```
switch# show interface fc3/13 , fc3/16
fc3/13 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:8d:00:05:30:00:97:9e
    Admin port mode is FX
    Port mode is F, FCID is 0x7b0300
    Port vsan is 1
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 12
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1856 frames input, 116632 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      1886 frames output, 887712 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 1 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.

fc3/16 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:90:00:05:30:00:97:9e
    Admin port mode is FX
    Port mode is F, FCID is 0x7d0100
    Port vsan is 3000
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 12
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 504 bits/sec, 63 bytes/sec, 0 frames/sec
    5 minutes output rate 520 bits/sec, 65 bytes/sec, 0 frames/sec
      47050 frames input, 10311824 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      62659 frames output, 10676988 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

Example 12-5 Display a Specific Interface

```
switch# show interface fc2/2
fc2/2 is trunking
    Port description is Trunk to Core-4
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:42:00:05:30:00:97:9e
    Peer port WWN is 20:cc:00:05:30:00:50:9e
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Transmit B2B Credit is 255
    Receive B2B Credit is 255
    Receive data field Size is 2112
    Beacon is turned off
    Belongs to port-channel 2
    Trunk vsans (admin allowed and active) (1,100,3000)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 (100,3000)
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
      2214834 frames input, 98673588 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      2262415 frames output, 343158368 bytes
        0 discards, 0 errors
      1 input OLS, 1 LRR, 1 NOS, 0 loop inits
      2 output OLS, 1 LRR, 0 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

Example 12-6 Displays Port Description

```
switch# show interface description
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
fc3/1              test intest
fc3/2              --
fc3/3              --
fc3/4              TE port
fc3/5              --
fc3/6              --
fc3/10             Next hop switch 5
fc3/11             --
fc3/12             --
fc3/16             --
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
port-channel 1     --
port-channel 5     --
port-channel 6     --
```

Example 12-7 Display Interface Information in a Brief Format

```
switch# show interface brief

--------------------------------------------------------------------------------
```

```
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper    Port
                  Mode   Trunk                          Mode  Speed   Channel
                         Mode                                 (Gbps)
--------------------------------------------------------------------------------
fc1/1      1      E      on      trunking         swl    TE      2     1
fc1/2      1      E      on      trunking         swl    TE      2     1
fc1/3      1      auto   on      SFPAbsent       --     --            --
fc1/4      1      auto   on      SFPAbsent       --     --            --
fc1/5      3000   auto   on      up               swl    F       2     --
...
fc2/2      1      E      on      trunking         swl    TE      2     2
fc2/3      1      auto   on      down             c1610  --            --
fc2/4      1      auto   on      down             c1590  --            --
fc2/5      3000   auto   on      notConnected     lwcr   --            --
fc2/6      1      auto   on      SFPAbsent       --     --            --
...
fc3/16     3000   FX     --      up               swl    F       2     --
fc3/17     1      FX     --      SFPAbsent       --     --            --
...
--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed        MTU
--------------------------------------------------------------------------------
GigabitEthernet4/1      SFPAbsent --                auto         1500
...
GigabitEthernet4/6      down       10.1.1.2/8       auto         3000
GigabitEthernet4/7      down       10.1.1.27/24     auto         1500
GigabitEthernet4/8      down       --               auto         1500

--------------------------------------------------------------------------------
Interface           Status              Oper Mode           Oper Speed
                                                            (Gbps)
--------------------------------------------------------------------------------
iscsi4/1            down                --
...
--------------------------------------------------------------------------------
Interface          Status                           Speed
                                                    (Gbps)
--------------------------------------------------------------------------------
sup-fc0            up                               1

--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed        MTU
--------------------------------------------------------------------------------
mgmt0                   up         172.19.48.96/25   100 Mbps     1500

--------------------------------------------------------------------------------
Interface               Vsan     Admin    Status            Oper   Oper
                                 Trunk                      Mode   Speed
                                 Mode                              (Gbps)
--------------------------------------------------------------------------------
port-channel 1          1        on       trunking          TE     4
port-channel 2          1        on       trunking          TE     4

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status       Oper  Profile  Port-channel
                  Mode   Trunk                Mode
                         Mode
--------------------------------------------------------------------------------
fcip10     1      auto   on      notConnected --     10       --
```

Example 12-8 Display Interface Counters

```
switch# show interface counters
```

```
fc3/1
    5 minutes input rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
    5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
    3502 frames input, 268400 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    3505 frames output, 198888 bytes
      0 discards
    1 input OLS, 1 LRR, 1 NOS, 0 loop inits
    2 output OLS, 1 LRR, 1 NOS, 0 loop inits
    1 link failures, 1 sync losses, 1 signal losses
.
.
.
fc9/8
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
sup-fc0
    114000 packets input, 11585632 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    113997 packets output, 10969672 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors

mgmt0
    31557 packets input, 2230860 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    26618 packets output, 16824342 bytes, 0 underruns
      0 output errors, 0 collisions, 7 fifo
      0 carrier errors

vsan1
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 1
```

```
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
```

Note Interfaces 9/8 and 9/9 are not trunking ports and display class 2, 3, and F information as well.

Example 12-9 Display Interface Counters in Brief Format

```
switch# show interface counters brief

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------ ------------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
-------------------------------------------------------------------------------
fc3/1              0        3871                  0        3874
fc3/2              0        3902                  0        4232
fc3/3              0        3901                  0        4138
fc3/4              0        3895                  0        3894
fc3/5              0        3890                  0        3897
fc9/8              0        0                     0        0
fc9/9              0        5                     0        4
fc9/10             0        4186                  0        4182
fc9/11             0        4331                  0        4315

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------ ------------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
-------------------------------------------------------------------------------
port-channel 1     0        0                     0        0
port-channel 2     0        3946                  0        3946
```

Note The **show interface transceiver** command can only be issued on a switch in the Cisco MDS 9100 Series if the SFP is present (see Example 12-10).

Example 12-10 Display Transceiver Information

```
switch# show interface transceiver
fc1/1 SFP is present
    name is CISCO-AGILENT
    part number is QFBR-5796L
    revision is
    serial number is A00162193
    fc-transmitter type is short wave laser
    cisco extended id is unknown (0x0)
...
fc1/9 SFP is present
    name is FINISAR CORP.
    part number is FTRJ-1319-7D-CSC
    revision is
    serial number is H11A6ER
    fc-transmitter type is long wave laser cost reduced
    cisco extended id is unknown (0x0)
...
```

Example 12-11 displays the entire running configuration with information for all interfaces. The interfaces have multiple entries in the configuration files to ensure that the interface configuration commands execute in the correct order when the switch reloads.

Example 12-11 Display the Running Configuration for All Interfaces

```
switch# show running-config
...
interface fc9/1
  switchport speed 2000
...
interface fc9/1
  switchport mode E
...
interface fc9/1
  channel-group 11 force
  no shutdown
```

Example 12-12 displays the running configuration information for a specified interface. The interface configuration commands are grouped together

Example 12-12 Display the Running Configuration for a Specified Interface

```
switch# show running-config interface fc1/1
interface fc9/1
  switchport speed 2000
  switchport mode E
  channel-group 11 force
  no shutdown
```

Example 12-13 displays the running configuration after the **system default switchport mode F** command is executed. Example 12-14 displays the running configuration after two interfaces are individually configured for mode FL.

Example 12-13 Display the Running Configuration After the System Default Switchport Mode F Command is Executed

```
switch# show running-config
version 3.1(3)
system default switchport mode F
```

```
interface fc4/1
interface fc4/2
interface fc4/3
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
interface fc4/8
interface fc4/9
interface fc4/10
```

Example 12-14 Display the Running Configuration After Two Interfaces Are Individually Configured for Mode FL

```
switch# show running-config
version 3.1(3)
system default switchport mode F
interface fc4/1
  switchport mode FL
interface fc4/2
interface fc4/3
  switchport mode FL
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
interface fc4/8
interface fc4/9
interface fc4/1
```

Example 12-15 displays interface information in a brief format after the **system default switchport mode F** command is executed. Example 12-16 displays interface information in a brief format after two interfaces are individually configured for mode FL.

Example 12-15 Display Interface Information in a Brief Format After the System Default Switchport Mode F Command is Executed

```
switch# sh int brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                          (Gbps)
--------------------------------------------------------------------------------
fc4/1      1      F      --      notConnected    swl    --           --
fc4/2      1      F      --      notConnected    swl    --           --
fc4/3      1      F      --      notConnected    swl    --           --
fc4/4      1      F      --      notConnected    swl    --           --
fc4/5      1      F      --      sfpAbsent       --     --           --
fc4/6      1      F      --      sfpAbsent       --     --           --
fc4/7      1      F      --      sfpAbsent       --     --           --
fc4/8      1      F      --      sfpAbsent       --     --           --
fc4/9      1      F      --      sfpAbsent       --     --           --
```

Example 12-16 Display Interface Information in a Brief Format After Two Interfaces Are Individually Configured for Mode FL

```
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
```

```
                          Mode                                    (Gbps)
--------------------------------------------------------------------------------
fc4/1       1       FL      --      notConnected     swl     --          --
fc4/2       1       F       --      notConnected     swl     --          --
fc4/3       1       FL      --      notConnected     swl     --          --
fc4/4       1       F       --      notConnected     swl     --          --
fc4/5       1       F       --      sfpAbsent        --      --          --
fc4/6       1       F       --      sfpAbsent        --      --          --
fc4/7       1       F       --      sfpAbsent        --      --          --
fc4/8       1       F       --      sfpAbsent        --      --          --
fc4/9       1       F       --      sfpAbsent        --      --          --
fc4/10      1       F       --      sfpAbsent        --      --          --
```

TL Ports for Private Loops

Private loops require setting the interface mode to TL. This section describes TL ports and includes the following sections:

About TL Ports

TL port mode is not supported on the following:

- Generation 2 switching module interfaces
- Cisco MDS 9124 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop.

The legacy devices are used in Fibre Channel networks and devices outside the loop may need to communicate with them.The communication functionality is provided through TL ports. See the "About Interface Modes" section on page 12-3.

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxied to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.

- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

Table 12-6 lists the TL port translations supported in Cisco MDS 9000 Family switches. Figure 12-4 shows examples of TL port translation support.

Table 12-6 Supported TL Port Translations

Translation from	Translation to	Example
Private initiator	Private target	From I1 to T1 or vice versa
Private initiator	Public target — N port	From I1 to T2 or vice versa
Private initiator	Public target — NL port	From I4 to T3 or vice versa
Public initiator — N port	Private target	From I2 to T1 or vice versa
Public initiator — NL port	Private target	From I3 to T1 or vice versa

Figure 12-4 TL Port Translation Support Examples

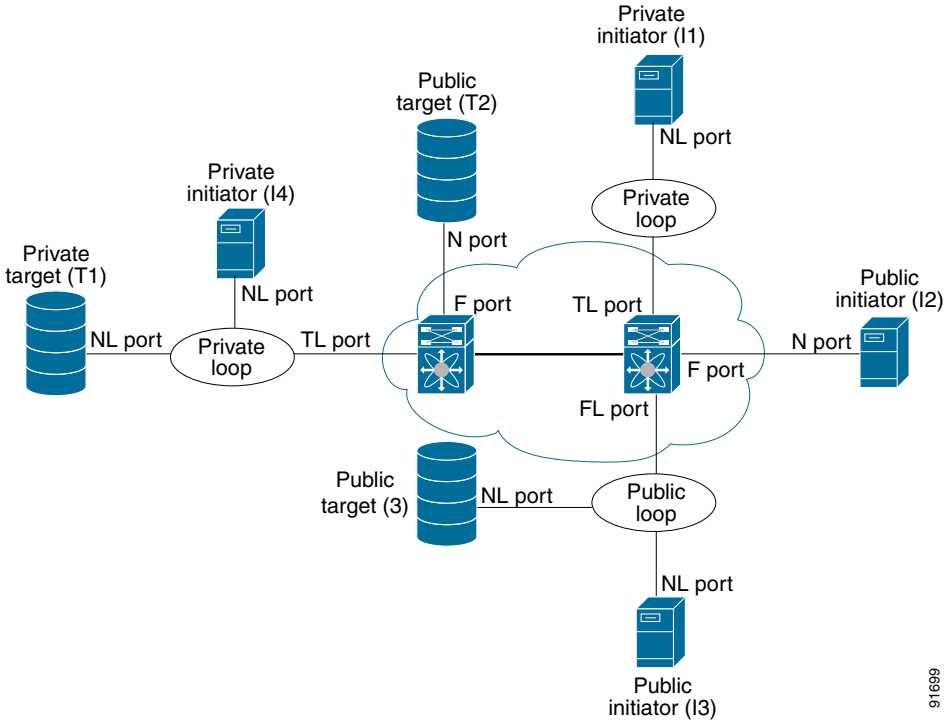

About TL Port ALPA Caches

Although TL ports cannot be automatically configured, you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco SAN-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco SAN-OS software discards an inactive cache entry (if available) to make space for the new entry. See the "TL Port" section on page 12-5 for more information on TL ports.

Displaying TL Port Information

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop.

The legacy devices are used in Fibre Channel networks and devices outside the loop may need to communicate with them.The communication functionality is provided through TL ports.

Use the **switchport mode** command to configure a TL port (see the "Configuring Interface Modes" section on page 12-13).

The **show tlport** command displays the TL port interface configurations. This command provides a list of all TL ports configured in a switch and shows the associated VSAN, the FC ID for the port (only domain and area are valid), and the current operational state of the TL port (up or initializing). See Example 12-17 through Example 12-20.

Example 12-17 Displays the TL Ports in All VSANs

```
switch# show tlport list
--------------------------------
Interface Vsan FC-ID    State
------------------------ ------
fc1/16    1    0x420000 Init
fc2/26    1    0x150000 Up
```

TL ports allow a private device (devices that physically reside on the loop) to see a fabric device and vice-versa by proxying fabric devices on the loop. Fabric devices are proxied by allocating each fabric device an ALPA on this loop.

In addition to these proxied devices, other virtual devices (local or remote domain controller addresses) are also allocated ALPAs on the loop. A switch reserves the ALPA for its own communication with private devices, and the switch acts as a SCSI initiator.

The first column in the output of the **show tlport interface** command is the ALPA identity of the device on the loop. The columns that follow include the port WWNs, the node WWNs for each device, the device as a SCSI initiator or target, and the real FC ID of the device.

Example 12-18 Displays the Detailed Information for a Specific TL Port

```
switch# show tlport interface fc1/16 all
fc1/16 is up, vsan 1, FCID 0x420000
-------------------------------------------------------------------------------
alpa pWWN                   nWWN                   SCSI Type Device  FC-ID
-------------------------------------------------------------------------------
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Proxied 0xfffc42
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target    Private 0x420073
0xef 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Switch  0x0000ef
```

Example 12-19 Displays TL Port Information for Private Devices

```
switch# show tlport interface fc 1/16 private
fc1/16 is up, vsan 1, FCID 0x420000
--------------------------------------------------------------------------
alpa pWWN                    nWWN                    SCSI Type FC-ID
--------------------------------------------------------------------------
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target    0x420073
0x74 22:00:00:20:37:38:d3:de 20:00:00:20:37:38:d3:de Target    0x420074
```

Example 12-20 Displays TL Port Information for Proxied Devices

```
switch# show tlport interface fc 1/16 proxied
fc1/16 is up, vsan 1, FCID 0x420000
--------------------------------------------------------------------------
alpa pWWN                    nWWN                    SCSI Type FC-ID
--------------------------------------------------------------------------
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator 0xfffc42
0x02 21:00:00:e0:8b:01:95:e7 20:00:00:e0:8b:01:95:e7 Initiator 0x420100
```

Manually Inserting Entries into ALPA Cache

To manually insert entries into the ALPA cache, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tlport alpa-cache interface fc1/2 pwwn 22:00:00:20:37:46:09:bd alpa 0x02`	Configures manual entries into the ALPA cache.
Step 3	`switch(config)# tlport alpa-cache interface fc1/3 pwwn 22:00:00:20:37:46:09:bd`	Removes this entry from the ALPA cache.

Displaying the ALPA Cache Contents

The **show tlport alpa-cache** command displays the contents of the ALPA cache.

```
switch# show tlport alpa-cache
-----------------------------------------------------
alpa              pWWN               Interface
-----------------------------------------------------
0x02  22:00:00:20:37:46:09:bd     fc1/2
0x04  23:00:00:20:37:46:09:bd     fc1/2
```

The first entry indicates that if a device with a pWWN of 22:00:00:20:37:46:09:bd is exported on TL port fc1/2, then the pWWN is allocated an alpa 0x02 (if available).

Clearing the ALPA Cache

The **clear tlport alpa-cache** command clears the entire content of the ALPA cache.

Buffer Credits

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination. This section describes the different buffer credits available on the Cisco MDS Family switches and includes the following topics:

About Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow control mechanism to ensure that FC switches do not run out of buffers, because switches must not drop frames. BB_credits are negotiated on a per-hop basis.

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

Note The receive BB_credit values depend on the module type and the port mode, as follows:

- For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
- For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.
- For Generation 2 switching modules, see the "Buffer Pools" section on page 14-8.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Configuring Buffer-to-Buffer Credits

To configure BB_credits for a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport fcrxbbcredit default`	Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities.
	`switch(config-if)# switchport fcrxbbcredit 5`	Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 255.
	`switch(config-if)# switchport fcrxbbcredit 5 mode E`	Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 255.
	`switch(config-if)# switchport fcrxbbcredit 5 mode Fx`	Assigns this value if the port is operating in F or FL mode. The range to assign BB_credits is between 1 and 255.

Step 4	`switch(config-if# do show int fc1/1`
	`fc1/1 is up`
	`...`
	`16 receive B2B credit remaining`
	`3 transmit B2B credit remaining`
		Displays the receive and transmit BB_credit along with other pertinent interface information for this interface.
		Note The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow.

About Performance Buffers

Note Performance buffers are not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCeter.

Regardless of the configured receive BB_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB_credit value.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

Configuring Performance Buffers

To configure performance buffers for a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport fcrxbbcredit performance-buffers 45`	Assigns a performance buffer of 45 to the selected interface. The value ranges from 1 to 145.
	`switch(config-if)# switchport fcrxbbcredit performance-buffers default`	Reverts to the factory default of using the built-in algorithm.

Note Use the **show interface bbcredit** command to display performance buffer values and other BB_credit information.

About Extended BB_credits

You can use the extended BB_credits flow control mechanism in addition to BB_credits for long haul links.

This section includes the following topics:

Extended BB_credits on Generation 1 Switching Modules

The BB_credits feature allows you to configure up to 255 receive buffers on Generation 1 switching modules. To facilitate BB_credits for long haul links, you can configure up to 3,500 receive BB_credits on a Fibre Channel port on a Generation 1 switching module.

To use this feature on Generation 1 switching modules, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license (see Chapter 3, “Obtaining and Installing Licenses”).
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see Figure 12-5).

Figure 12-5 Port Group Support for the Extended BB_Credits Feature

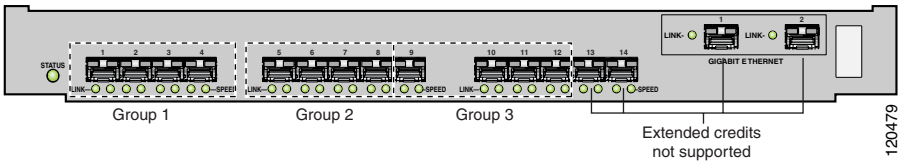

The port groups that support extended credit configurations are as follows.

- Any one port in ports 1 to 4 (identified as Group 1 in Figure 12-1).
- Any one port in ports 5 to 8 (identified as Group 2 in Figure 12-1).
- Any one port in ports 9 to 12 (identified as Group 3 in Figure 12-1).

Note The last two Fibre Channel ports (port 13 and port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature (see Figure 12-1).

- Explicitly enable this feature in the required Cisco MDS switch.
- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
  - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.

    Note The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

  - If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.
- Be aware that changing the BB_credit value results in the port being disabled and then reenabled.
- Disable (explicitly) this feature if you need to nondisruptively downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.

Note The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

Extended BB_credits on Generation 2 Switching Modules

To use this feature on Generation 2 switching modules, you must meet the following requirements:

- Obtain the Enterprise package (ENTERPRISE_PKG) license (see Chapter 3, "Obtaining and Installing Licenses").
- Configure this feature in any port on a Generation 2 switch module. See the Extended BB_Credits, page 14-15 for more information on extended BB_credits on Generation 2 switching modules.

Note Extended BB_credits are not supported on the Cisco MDS 9124 Fabric Switch.

Configuring Extended BB_credits

To configure extended BB_credits for a MDS-14/2 interface, for a Generation 2 switching module interface (not including the Cisco MDS 9124 Fabric Switch), or for an interface in a Cisco MDS 9216i switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcrxbbcredit extended enable`	Enables the extended BB_credits feature.
	`switch(config)# no fcrxbbcredit extended enable`	Disables (default) the extended BB_credits feature.

	Command	Purpose
Step 3	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 4	`switch(config-if)# switchport fcrxbbcredit extended 1500`	Applies the extended BB_credit value of 1,500 credits to the selected interface. The valid range is from 256 to 3,500 credits.
	`switch(config-if)# no switchport fcrxbbcredit extended 1500`	Clears the configured extended BB_credit configuration for this port.

Step 5	`switch# do show interface fc3/2`
	`fc3/2 is trunking`
	`Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)`
	`Port WWN is 20:82:00:05:30:00:2a:1e`
	`Peer port WWN is 20:42:00:0b:46:79:f1:80`
	`Admin port mode is auto, trunk mode is on`
	`Port mode is TE`
	`Port vsan is 1`
	`Speed is 2 Gbps`
	`Transmit B2B Credit is 255`
	`Receive B2B Credit is 1500`
	`Receive data field Size is 2112`
	`...`
		Displays the receive and transmit BB_credit values along with other pertinent interface information for this interface if the interface is in the up state.
		Note The receive BB_credit value reflects the extended BB_credit configuration, if applicable.

Displaying BB_Credit Information

To display the BB_credit information, use the **show interface bbcredit** command (see Example 12-21 and Example 12-22).

Example 12-21 Displays BB_credit Information

```
switch# show interface bbcredit
fc2/1 is down (SFP not present)
...
fc2/17 is trunking
    Transmit B2B Credit is 255
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 375
      12 receive B2B credit remaining
      255 transmit B2B credit remaining
fc2/18 is down (SFP not present)
fc2/19 is down (SFP not present)
fc2/20 is down (SFP not present)
fc2/21 is down (Link failure or not-connected)
...
fc2/31 is up
    Transmit B2B Credit is 0
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 48
      12 receive B2B credit remaining
      0 transmit B2B credit remaining
fc2/32 is down (Link failure or not-connected)
```

Example 12-22 Displays BB_credit Information for a Specified Fibre Channel Interface

```
switch# show interface fc2/31 bbcredit
fc2/31 is up
    Transmit B2B Credit is 0
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 48
      12 receive B2B credit remaining
      0 transmit B2B credit remaining
```

Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection on the mgmt0 interface, you must configure either the IP version 4 (IPv4) parameters (IP address, subnet mask, and default gateway) or the IP version 6 (IPv6) parameters so that the switch is reachable.

This section describes the management interfaces and includes the following topics:

About Management Interfaces

Before you begin to configure the management interface manually, obtain the switch's IPv4 address and subnet mask, or the IPv6 address.

The management port (mgmt0) is autosensing and operates in full duplex mode at a speed of 10/100/1000 Mbps. Autosensing supports both the speed and the duplex mode. On a Supervisor-1 module, the default speed is 100 Mbps and the default duplex mode is auto. On a Supervisor-2 module, the default speed is auto and the default duplex mode is auto.

Note You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Configuring Management Interfaces

To configure the mgmt0 Ethernet interface to connect over IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Selects the management Ethernet interface on the switch and enters interface configuration submode.
Step 3	`switch(config-if)# ip address 10.16.1.2 255.255.255.0`	Configures the IPv4 address and IPv4 subnet mask.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

	Command	Purpose
Step 5	`switch(config-if)# exit` `switch(config)#`	Returns to configuration mode.
Step 6	`switch(config)# ip default-gateway 1.1.1.4`	Configures the default gateway IPv4 address.
Step 7	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 8	`switch# copy running-config startup-config`	(Optional) Saves your configuration changes to the file system. **Note** If you wish to save your configuration, you can issue this command at any time.

To configure the mgmt0 Ethernet interface to connect over IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Selects the management Ethernet interface on the switch and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 enable`	Enables IPv6 and assigns a link-local address on the interface.
Step 4	`switch(config-if)# ipv6 address ipv6 address 2001:0db8:800:200c::417a/64`	Specifies an IPv6 unicast address and prefix length on the interface.
Step 5	`switch(config-if)# no shutdown`	Enables the interface.
Step 6	`switch(config-if)# end` `switch#`	Returns to EXEC mode.
Step 7	`switch# copy running-config startup-config`	(Optional) Saves your configuration changes to the file system. **Note** If you wish to save your configuration, you can issue this command at any time.

Displaying Management Interface Configuration

To display the management interface configuration, use the **show interface mgmt 0** command.

```
switch# show interface mgmt 0
mgmt0 is up
    Hardware is FastEthernet
    Address is 000c.30d9.fdbc
    Internet address is 10.16.1.2/24
    MTU 1500 bytes, BW 100 Mbps full Duplex
    26388 packets input, 6101647 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    10247 packets output, 2389196 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexisting VSANs.

This section describes VSAN interfaces and includes the following topics:

About VSAN Interfaces

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.

Tip After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature (see Chapter 43, "Configuring IP Services").

Creating VSAN Interfaces

To create a VSAN interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 2` `switch(config-if)#`	Configures a VSAN with the ID 2.
Step 3	`switch(config-if)# no shutdown`	Enables the VSAN interface.

Displaying VSAN Interface Information

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

```
switch# show interface vsan 2
vsan2 is up, line protocol is up
    WWPN is 10:00:00:05:30:00:59:1f, FCID is 0xb90100
    Internet address is 10.1.1.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
```

Default Settings

Table 12-7 lists the default settings for interface parameters.

Table 12-7 Default Interface Parameters

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

CHAPTER 13

Configuring N Port Virtualization

N Port virtualization (NPV) reduces the number of Fibre Channel domain IDs in SANs. Switches operating in the NPV mode do not join a fabric; rather, they pass traffic between NPV core switch links and end devices, which eliminates the domain IDs for these edge switches.

NPV is supported by the following Cisco MDS 9000 switches only:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

Note NPV is available on these switches only while in NPV mode; if in switch mode, NPV is not available.

This chapter includes the following sections:

About NPV

Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to core devices. However, as the number of ports in the fabric increases, the number of switches deployed also increases, and you can end up with a dramatic increase in the number of domain IDs (the maximum number supported is 239). This challenge becomes even more difficult when additional blade chassis are deployed in Fibre Channel networks.

NPV addresses the increase in the number of domain IDs needed to deploy a large number of the ports by making a fabric or module switch appear as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric or blade switch. NPV aggregates multiple locally connected N ports into one or more external NP links, which shares the domain ID of the NPV core switch among multiple NPV switches (see Figure 13-1). NPV also allows multiple devices to attach to the same port on the NPV core switch, thereby reducing the need for more ports on the core.

Figure 13-1 Cisco NPV Fabric Configuration

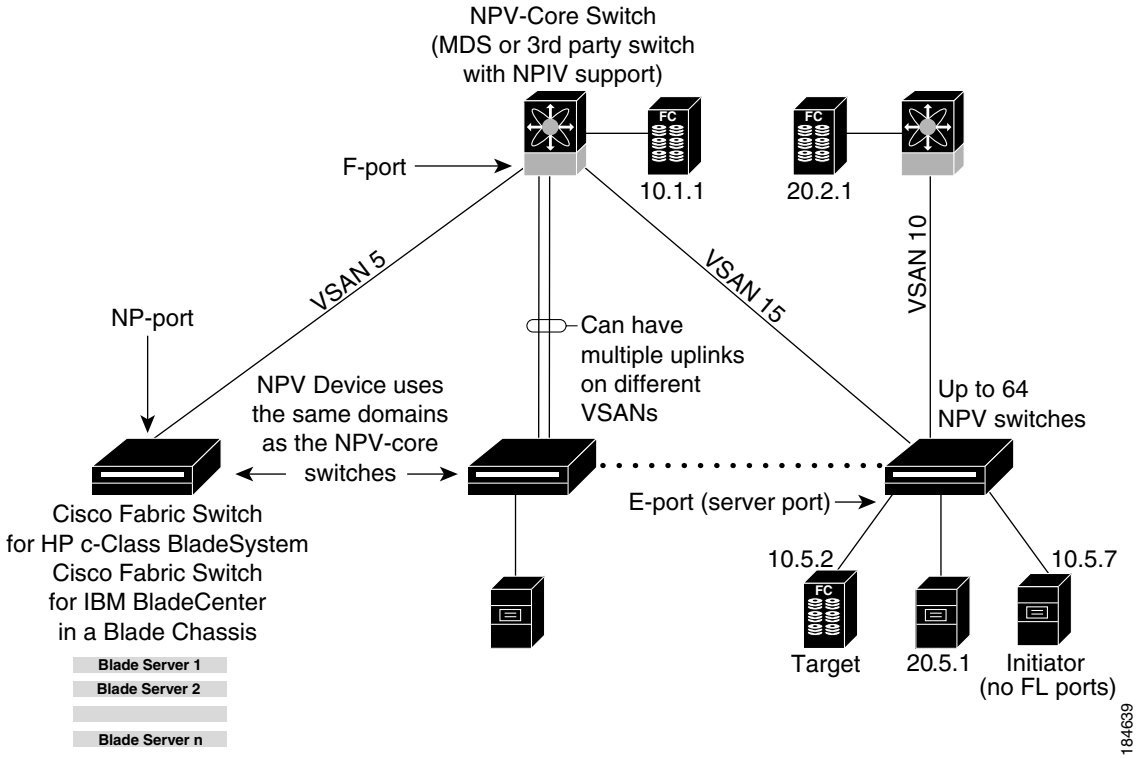

While NPV is similar to N port identifier virtualization (NPIV), it does not offer exactly the same functionality. NPIV provides a means to assign multiple FC IDs to a single N port, and allows multiple applications on the N port to use different identifiers. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of NPIV to get multiple FCIDs allocated from the core switch on the NP port.

Figure 13-2 shows a more granular view of an NPV configuration at the interface level.

Figure 13-2 Cisco NPV Configuration–Interface View

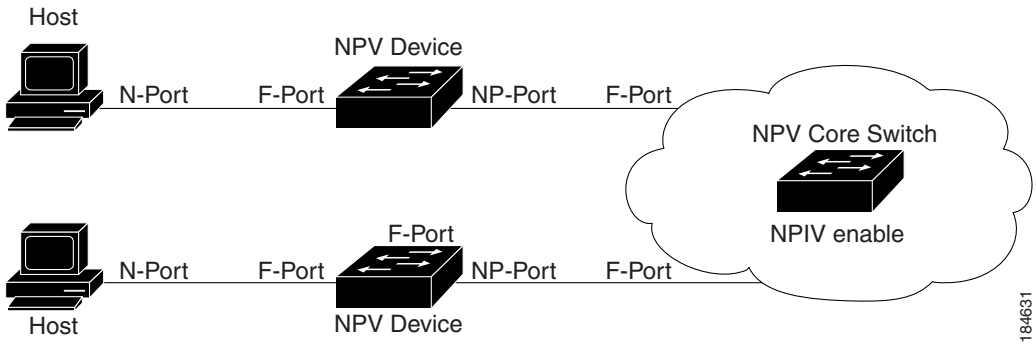

NPV Mode

A switch is in NPV mode after a user has enabled NPV and the switch has successfully rebooted. NPV mode applies to an entire switch. All end devices connected to a switch that is in NPV mode must log in as an N port to use this feature (loop-attached devices are not supported). All links from the edge switches (in NPV mode) to the NPV core switches are established as NP ports (not E ports), which are used for typical interswitch links. NPIV is used by the switches in NPV mode to log in to multiple end devices that share a link to the NPV core switch.

Note In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink to the core from the NPV device. For traffic beyond the NPV device, core switches will enforce in-order delivery if needed and/or configured.

After entering NPV mode, only the following commands are available:

```
aaa             Configure aaa functions
arp             [no] remove an entry from the ARP cache
banner          Configure banner message
boot            Configure boot variables
callhome        Enter the callhome configuration mode
cli             CLI configuration commands
clock           Configure time-of-day clock
do              EXEC command
end             Exit from configure mode
exit            Exit from configure mode
fcanalyzer      Configure cisco fabric analyzer
fcrxbbcredit    Enable extended rx b2b credit configuration
fips            Enable/Disable FIPS mode
hw-module       Enable/Disable OBFL information
interface       Select an interface to configure
ip              Configure IP features
ipv6            Configure IPv6 features
line            Configure a terminal line
logging         Modify message logging facilities
no              Negate a command or set its defaults
npv             Config commands for FC N_port Virtualizer
ntp             NTP Configuration
port-track      Configure Switch port track config
power           Configure power supply
poweroff        Poweroff a module in the switch
radius          Configure RADIUS configuration
radius-server   Configure RADIUS related parameters
```

```
rate-mode       Configure rate mode oversubscription limit
rmon            Remote Monitoring
role            Configure roles
snmp-server     Configure snmp server
ssh             Configure SSH parameters
switchname      Configure system's network name
system          System config command
tacacs+         Enable tacacs+
telnet          Enable telnet
username        Configure user information.
wwn             Set secondary base MAC addr and range for additional WWNs
```

NP Ports

An *NP port* (proxy N port) is a port on a device that is in NPV mode and connected to the NPV core switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

NP Links

An *NP link* is basically an NPIV uplink to a specific end device. NP links are established when the uplink to the NPV core switch comes up; the links are terminated when the uplink goes down. Once the uplink is established, the NPV switch performs an internal FLOGI to the NPV core switch, and then (if the FLOGI is successful) registers itself with the NPV core switch's name server. Subsequent FLOGIs from end devices in this NP link are converted to FDISCs. For more details refer to the "Internal FLOGI Parameters" section on page 13-4.

Server links are uniformly distributed across the NP links. All the end devices behind a server link will be mapped to only one NP link.

Internal FLOGI Parameters

When an NP port comes up, the NPV device first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fWWN (fabric port WWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based sWWN (switch WWN) of the NPV device used as nWWN (node WWN) in the internal FLOGI.

After completing its FLOGI request, the NPV device registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc1/4) of the NP port is embedded in the symbolic port name in the name server registration of the NPV device itself.
- The IP address of the NPV device is registered as the IP address in the name server registration of the NPV device.

Note The BB_SCN of internal FLOGIs on NP ports is always set to zero. The BB_SCN is supported at the F-port of the NPV device.

Figure 13-3 shows the internal FLOGI flows between an NPV core switch and an NPV device.

Figure 13-3 Internal FLOGI Flows

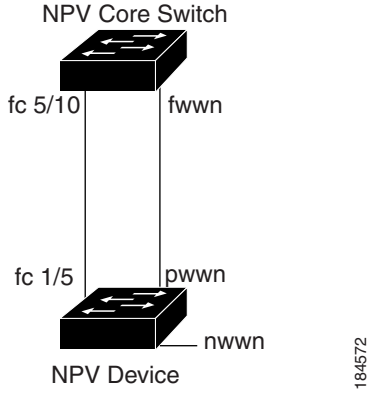

Table 13-1 identifies the internal FLOGI parameters that appear in Figure 13-3.

Table 13-1 Internal FLOGI parameters

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the NPV device.
fWWN	The fWWN of the F port on the NPV core switch.
symbolic port name	The switch name and NP port interface string. **Note** If there is no switch name available, then the output will simply read "switch." For example, `switch: fc1/5`.
IP address	The IP address of the NPV device.
symbolic node name	The NPV switch name.

Although fWWN-based zoning is supported for NPV devices, it is not recommended because:

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

Default Port Numbers

Port numbers on NPV-enabled switches will vary depending on the switch model. For details about port numbers for NPV-eligible switches, see Chapter 4, "On-Demand Port Activation Licensing."

NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- NPV core switches must support NPIV.
- You can have up to 100 NPV devices.

- Nondisruptive upgrades are supported. See Chapter 7, "Software Images."
- Port tracking is supported. See Chapter 57, "Configuring Port Tracking."
- You can configure zoning for end devices that are connected to NPV devices using all available member types on the NPV core switch. If fWWN, sWWN, domain, or port-based zoning is used, then fWWN, sWWN or the domain/port of the NPV core switch should be used.
- Port security is supported on the NPV core switch for devices logged in via NPV.
- NPV uses a load balancing algorithm to automatically assign end devices in a VSAN to one of the NPV core switch links (in the same VSAN) upon initial login. If there are multiple NPV core switch links in the same VSAN, then you cannot assign a specific one to an end device.
- Both servers and targets can be connected to an NPV device.
- Remote SPAN is not supported.
- Local switching is not supported; all traffic is switched using the NPV core switch.
- NPV devices can connect to multiple NPV core switches. In other words, different NP ports can be connected to different NPV core switches.
- NPV supports NPIV-capable module servers (nested NPIV).
- Only F, NP, and SD ports are supported in NPV mode.

Configuring NPV

When you enable NPV, your system configuration is erased and the system is rebooted with NPV mode enabled.

Note We recommend that you save your current configuration either bootflash or a TFTP server before NPV (if the configuration is required for later use). Use the following commands to save either your non-NPV or NPV configuration:

```
switch# copy running bootflash:filename
```

The configuration can be reapplied later using the following command:

```
switch# copy bootflash:filename running-config
```

SUMMARY STEPS

5. Enable NPIV on the NPV core switch.Enable NPV on the NPV device.
6. Configure the interfaces connected to the NPV core switch as NP ports.
7. Configure NPV link as an F port on the NPV core switch.
8. Configure server link on the NPV device.

Note On the 91x4 platform, before you upgrade to 3.2(2c) or downgrade from 3.2(2c), shut the F-ports connected to NPIV capable hosts, and then disable the NPIV feature. After the upgrade or downgrade is complete, enable the NPIV feature and up the F-ports.

On the 91x4 platform, before you downgrade from 3.2(2c) to prior versions, shut the F-port, enable and disable the FC domain persistency for that VSAN and then up the F-port. Use the following steps to configure NPV.

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	On the NPV core switch, enters configuration mode.
Step 2	`switch(config)# npiv enable` `switch (config)#`	Enables NPIV mode on the NPV core switch.
	`switch (config)# no npiv enable`	Disables NPIV mode on the NPV core switch.
Step 3	`switch(config)# interface fc2/1` `switch(config-if)# switchport mode F`	Configure the NPIV core switch port as an F port.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.
Step 4	`switch(config)# npv enable`	Enables NPV mode on a NPV device (module, Cisco MDS 9124 or Cisco MDS 9134 Fabric Switch). The module or switch is rebooted, and when it comes back up, is in NPV mode. **Note** A write-erase is performed during the reboot.
Step 5	`switch(config)# interface fc1/1` `switch(config-if)# switchport mode NP`	On the NPV device, select the interfaces that will be connected to the aggregator switch and configure them as NP ports.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.
Step 6	`switch(config-if)# exit`	Exits interface mode for the port.
Step 7	`switch(config)# interface fc1/2` `switch(config-if)# switchport mode F`	Selects the remaining interfaces on the NPV-enabled device and configures them as F ports.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.
Step 8	`switch(config-npv)# no npv enable` `switch(config)#`	Terminates session and disables NPV mode, which results in a reload of the NPV device.

Multiple VSAN Support

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs at the NPV-enabled switch. The correct uplink must be selected based on the VSAN(s) that the uplink can carry.

DPVM Configuration

When NPV is enabled, the following requirements must be met before you configure DPVM on the NPV core switch:

- You must explicitly configure the WWN of the internal FLOGI in DPVM.

- If DPVM is configured on the NPV core switch for an end device that is connected to the NPV device, then that end device must be configured to be in the same VSAN. Logins from a device connected to an NPV device will fail if the device is configured to be in a different VSAN. To avoid VSAN mismatches, ensure that the internal FLOGI VSAN matches the port VSAN of the NP port.
- The first login from an NP port determines the VSAN of that port. If DPVM is configured for this first login–which is the internal login of the NPV device–then the NPV core switch's VSAN F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see Chapter 21, "Creating Dynamic VSANs."

NPV and Port Security

Port security is enabled on the NPV core switch on a per interface basis. To enable port security on the NPV core switch for devices logging in via NPV, you must adhere to the following requirements:

- The internal FLOGI must be in the port security database; in this way, the port on the NPV core switch will allow communications/links.
- All the end device pWWNs must also be in the port security database.

Once these requirements are met, you can enable port security as you would in any other context. For details about enabling port security, see Chapter 38, "Configuring Port Security."

Verifying NPV

To view all the NPV devices in all the VSANs that the aggregator switch belongs to, enter the **show fcns database** command.

```
switch# show fcns database

VSAN 1:
--------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init

Total number of entries = 4
```

For additional details (such as IP addresses, switch names, interface names) about the NPV devices you see in the **show fcns database** output, enter the **show fcns database detail** command.

```
switch# show fcns database detail

------------------------
VSAN:1 FCID:0x010000
------------------------
port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/1
symbolic-node-name :para-3
port-type :N
```

```
port-ip-addr :0.0.0.0
fabric-port-wwn :20:01:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)

------------------------
VSAN:1 FCID:0x010001
------------------------
port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/2
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:02:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
```

If you need to contact support, enter the **show tech-support NPV** command and save the output so that support can use it to troubleshoot, if necessary.

To display a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, enter the **show npv flogi-table** command.

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID            PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
fc1/19    1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc1/1
fc1/19    1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc1/1

Total number of flogi = 4.
```

To display the status of the different servers and external interfaces, enter the **show npv status** command.

```
switch# show npv status
npiv is enabled

External Interfaces:
====================
  Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
  Interface: fc1/2, VSAN: 3, FCID: 0x040000, State: Up

  Number of External Interfaces: 2

Server Interfaces:
==================
  Interface: fc1/7, VSAN: 2, NPIV: No, State: Up
  Interface: fc1/8, VSAN: 3, NPIV: No, State: Up

  Number of Server Interfaces: 2
```

CHAPTER 14

Configuring Generation 2 Switches and Modules

The Cisco MDS 9500 Series switches and Cisco MDS 9216A and Cisco MDS 9216i switches support a set of modules called Generation 2 modules. This chapter describes how to configure these modules, as well as Generation 2 Multilayer Fabric Switches.

This chapter includes the following sections:

About Generation 2 Modules and Switches

Table 14-1 identifies the modules supported by the Cisco MDS 9500 Series switches and Cisco MDS 9216A and Cisco MDS 9216i switches, as well as the Fabric switches:

Table 14-1 Generation 2 Fibre Channel Modules and Fabric Switches

Part Number	Product Name/Description
Module	
DS-X9148	48-port 4-Gbps Fibre Channel switching module
DS-X9134	34-port 4-Gbps Fibre Channel switching module
DS-X9124	24-port 4-Gbps Fibre Channel switching module
DS-X9304-18K9	18-port 4-Gbps Fibre Channel switching module with 4 GigabitEthernet ports
DS-X9112	12-port 4-Gbps Fibre Channel switching module
DS-X9704	4-port 10-Gbps Fibre Channel switching module
DS-X9530-SF2-K9	Supervisor-2 module (Cisco MDS 9500 Series switches only)
Switch	

Table 14-1 Generation 2 Fibre Channel Modules and Fabric Switches (continued)

Part Number	Product Name/Description
DS-C9134-K9	Cisco MDS 9134 Fabric switch 32-port 4-Gbps Fabric switch with 2 additional 10-Gbps ports
DS-C9124	Cisco MDS 9124 Fabric switch 24-port 4-Gbps Fabric switch
DS-C9222i-K9	Cisco MDS 9222i Multiservice Modular switch 18-port 4-Gbps switch with 4 GigabitEthernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family Switching and Services Modules

Note Generation 2 Fibre Channel switching modules are not supported on the Cisco MDS 9216 switch; however, they are supported by both the Supervisor-1 module and the Supervisor-2 module.

For detailed information about the installation and specifications for these modules and switches, refer to the hardware installation guide for your switch.

This section includes the following topics:

Port Groups

Each module or switch can have one or more ports in port groups that share common resources (such as bandwidth and buffer credits). Table 14-2 shows the port groups for the Generation 2 Fibre Channel switches and modules.

Table 14-2 Bandwidth and Port Groups for Generation 2 FC Modules and Fabric Switches

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group	Maximum Bandwidth Per Port
Module				
DS-X9148	Cisco 48-port 4-Gbps Fibre Channel module 48-port 4-Gbps Fibre Channel switching module[1]	12	12.8	4-Gbps

Table 14-2 Bandwidth and Port Groups for Generation 2 FC Modules and Fabric Switches

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group	Maximum Bandwidth Per Port
DS-X9124	Cisco 24-port 4-Gbps Fibre Channel module 24-port 4-Gbps Fibre Channel switching module	6	12.8	4
DS-X9304-18K9	Cisco 18-port Fibre Channel /4-port GigabitEthernet Multiservice (MSM-18/4) module 18-port 4-Gbps Fibre Channel switching module with 4 GigabitEthernet ports	6	12.8	4-Gbps
DS-X9112	Cisco 12-port 4-Gbps Fibre Channel module 12-port 4-Gbps Fibre Channel switching module	3	12.8	4-Gbps
DS-X9704	Cisco 4-port 10-Gbps Fibre Channel module 4-port 10-Gbps Fibre Channel switching module	1	10	10-Gbps
Switches				
DS-C9134-K9	Cisco MDS 9134 Fabric switch 32-port 4-Gbps Fabric switch	4	16	4-Gbps
	2-port 10-Gbps Fabric switch	1	10	10-Gbps

Table 14-2 Bandwidth and Port Groups for Generation 2 FC Modules and Fabric Switches

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group	Maximum Bandwidth Per Port
DS-C9124-K9	Cisco MDS 9124 Fabric switch 24-port 4-Gbps	4	16	4-Gbps
DS-C9222i-K9	Cisco MDS 9222i Multiservice Modular switch 18-port 4-Gbps	6	12.8	4-Gbps

1. By default, all ports in a 48-port 4-Gbps switching module operate in shared mode with administrative operating speed set to auto. All ports in a 48-port 4-Gbps switching module can operate in dedicated mode with a 1-Gbps operating speed. However, if you configure one or more ports to operate in 2-Gbps or 4-Gbps dedicated mode, some of the other ports in the module would have to operate in shared mode.

Note Port groups are defined by the hardware and consist of sequential ports. For example, ports 1 through 12, ports 13 through 24, ports 25 through 36, and ports 37 through 48 are the port groups on the 48-port 4-Gbps Fibre Channel switching modules.

Port Rate Modes

The *Port rate mode* configuration is used to determine the bandwidth allocation for ports in a port group. Two port rate modes are supported: Dedicated Mode and Shared Mode. In Generation 1 modules, port rate mode is not configurable by users; rather, it is determined implicitly based on the port mode and linecard type. In Generation 2 modules, port rate mode is user-configured.

Table 14-3 Port Rate Mode Support on Generation 2 Modules and Switches

Part Number	Product Name/ Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode
Modules			
DS-X9148	Cisco 48-port 4-Gbps Fibre Channel module 48-port 4-Gbps Fibre Channel switching module[1]	Yes	Yes
DS-X9124	Cisco 24-port 4-Gbps Fibre Channel module 24-port 4-Gbps Fibre Channel switching module	Yes	Yes

Table 14-3 Port Rate Mode Support on Generation 2 Modules and Switches (continued)

Part Number	Product Name/ Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode
DS-X9304-18K9	Cisco 18-port Fibre Channel /4-port GigabitEthernet Multiservice (MSM-18/4) module 18-port 4-Gbps Fibre Channel switching module with 4 GigabitEthernet ports	Yes	Yes
DS-X9112	12-port 4-Gbps Fibre Channel module 12-port 4-Gbps Fibre Channel switching module	Yes	No
DS-X9704	4-port 10-Gbps Fibre Channel module 4-port 10-Gbps Fibre Channel switching module	Yes	No
Switches			
DS-C9134-K9	Cisco MDS 9134 Fabric switch 32-port 4-Gbps Fabric switch	Yes	Yes
	2-port 10-Gbps Fabric switch	Yes	No
DS-C9124	Cisco MDS 9124 Fabric switch 24-port 4-Gbps Fabric switch[2]	Yes	No
DS-C9222i-K9	Cisco MDS 9222i Multiservice Modular switch 18-port 4-Gbps Fibre Channel switch with 4 GigabitEthernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family Switching and Services Modules	Yes	Yes

1. By default, all ports in a 48-port 4-Gbps switching module operate in shared mode with administrative operating speed set to auto. All ports in a 48-port 4-Gbps switching module can operate in dedicated mode with a 1-Gbps operating speed. However, if you configure one or more ports to operate in 2-Gbps or 4-Gbps dedicated mode, some of the other ports in the module would have to operate in shared mode.

2. By default, all ports in a 24-port 4-Gbps switching module operate in shared mode with administrative operating speed set to auto. All ports in a 24-port 4-Gbps switching module can operate in dedicated mode with a 2-Gbps operating speed. However, if you configure one or more ports to operate in 4-Gbps dedicated mode, some of the other ports in the module would have to operate in shared mode

Note Port rate modes are supported on all Generation 2 modules and fabric switches. Port rate modes are not supported on the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Dedicated Mode

When port rate mode is configured as dedicated, a port is allocated required fabric bandwidth and related resources to sustain line rate traffic at the maximum operating speed configured for the port. In this mode, ports do not use local buffering and all receive buffers are allocated from a global buffer pool (see the "Buffer Pools" section on page 14-8).

Table 14-4 show the amount of bandwidth reserved for a configured port speed on 4-Gbps switching modules.

Table 14-4 Bandwidth Reserved for the Port Speeds on 4-Gbps Switching Modules

Configured Speed	Reserved Bandwidth
Auto	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

Note 10-Gbps ports in auto mode only support auto speed mode at 10 Gbps.

Shared Mode

When port rate mode is configured as shared, multiple ports within a port group share data paths to the switch fabric so that fabric bandwidth and related resources are shared. Often, the available bandwidth to the switch fabric may be less than the negotiated operating speed of a port. Ports in this mode use local buffering for the BB_credit buffers.

All ports in switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, or 4-Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated mode with 1-Gbps, 2-Gbps or 4-Gbps operating speed.

Dynamic Bandwidth Management

On port switching modules where bandwidth is shared, the bandwidth available to each port within a port group can be configured based on the port rate mode and speed configurations. Within a port group, some ports can be configured in dedicated mode while others operate in shared mode.

Ports configured in dedicated mode are allocated the required bandwidth to sustain a line rate of traffic at the maximum configured operating speed, and ports configured in shared mode share the available remaining bandwidth within the port group. Fair allocation of bandwidth among a group of ports is determined, in part, by the rate mode and speed configurations. For example, if the set ports in a module are configured with the same rate mode and speed (such as 4 Gbps of shared bandwidth), then all the ports should have fair allocation of bandwidth and eventually, similar throughput. When you enable bandwidth fairness, you should notice a reduction in any disparity that may otherwise exist in similar configurations.

Bandwidth allocation among the shared mode ports is based on the operational speed of the ports. For example, if four ports operating at speeds 1 Gbps, 1 Gbps, 2 Gbps, and 4 Gbps share bandwidth of 8 Gbps, the ratio of allocation would be 1:1:2:4.

Note If dedicated ports are not using all of their allocated bandwidth, the unused bandwidth is made available for use by all ports configured for shared bandwidth mode.

Tip When migrating a host that supports up to 2-Gbps traffic (that is, not 4-Gbps with autosensing capabilities) to the 4-Gbps switching modules, use autosensing with a maximum bandwidth of 2-Gbps.

Note If you configure an interface for autosensing speed with a maximum bandwidth of 2 Gbps and want to change to the default of 4 Gbps, ensure that there are enough shared resources available to support the configuration on the module.

Out-of-Service Interfaces

On supported modules and fabric switches, you might need to allocate all the shared resources for one or more interfaces to another interface in the port group or module. You can take interfaces out of service to release shared resources that are needed for dedicated bandwidth. When an interface is taken out of service, all shared resources are released and made available to the other interface in the port group or module. These shared resources include bandwidth, rate mode, BB_credits, and extended BB_credits. All shared resource configurations are returned to their default values when the interface is brought back into service. Corresponding resources must be made available in order for the port to be successfully returned to service.

Caution If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces in the same port group.

Buffer Credit Allocation

This sections describe how buffer credits are allocated to switches and modules, and includes the following topics:

Buffer Pools

In the architecture of Generation 2 modules, receive buffers shared by a set of ports are called *buffer groups*. The receive buffer groups are organized into *global* and *local* buffer pools.

The receive buffers allocated from the global buffer pool to be shared by a port group are called a *global buffer pool*. Global receive buffer pools include the following buffer groups:

- Reserved internal buffers
- Allocated BB_credit buffers for each Fibre Channel interface (user configured or assigned by default)
- Common unallocated buffer pool for BB_credits, if any, to be used for additional BB_credits as needed
- Performance buffers (only used on 12-port 4-Gbps and 4-port 10-Gbps switching modules)

Figure 14-1 shows the allocation of BB_credit buffers on linecards (24-port and 48-port line cards).

Note In some modules, performance buffers are not supported.

Figure 14-1 Receive Buffers for Fibre Channel Ports in a Global Buffer Pool

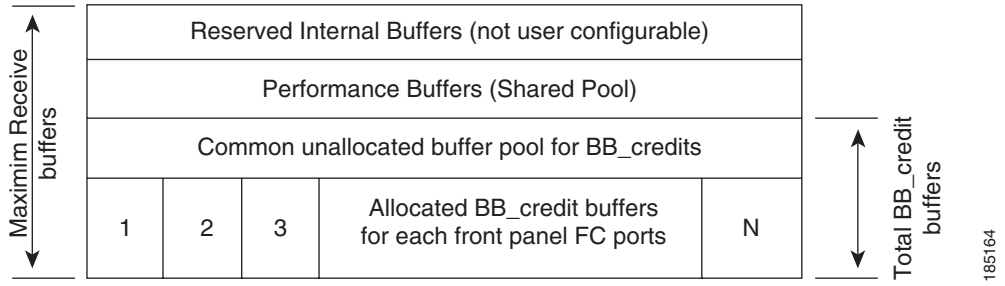

Figure 14-2 shows the default BB_credit buffer allocation model for 24-port 4-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 14-2 BB_Credit Buffer Allocation in 24-port 4-Gbps Switching Modules

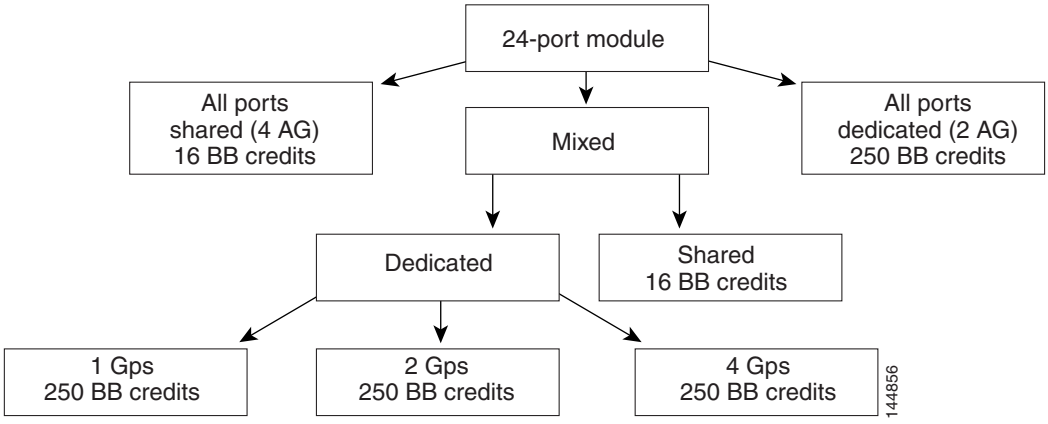

Note The default BB_credit buffer allocation is the same for all port speeds.

BB_Credit Buffers for Switching Modules

This section describes how buffer credits are allocated to Cisco MDS 9000 switching modules, and includes the following topics:

48-port 4-Gbps Fibre Channel Module BB_Credit Buffers

Table 14-5 lists the BB_credit buffer allocation for 48-port 4-Gbps Fibre Channel switching modules.

Table 14-5 48-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults			
		Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
		ISL[1]	Fx Port	ISL[1]	Fx Port
User configurable BB_credit buffers	6000	125	16	16	16

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 48-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 48-port 4-Gbps Fibre Channel switching module consists of 12 ports. The ports in shared rate mode have bandwidth oversubscription of 4:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 5:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 48-port 4-Gbps Fibre Channel switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)

- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 1-Gbps speed plus
  five ports put out-of-service (see Figure 14-3)

Figure 14-3 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

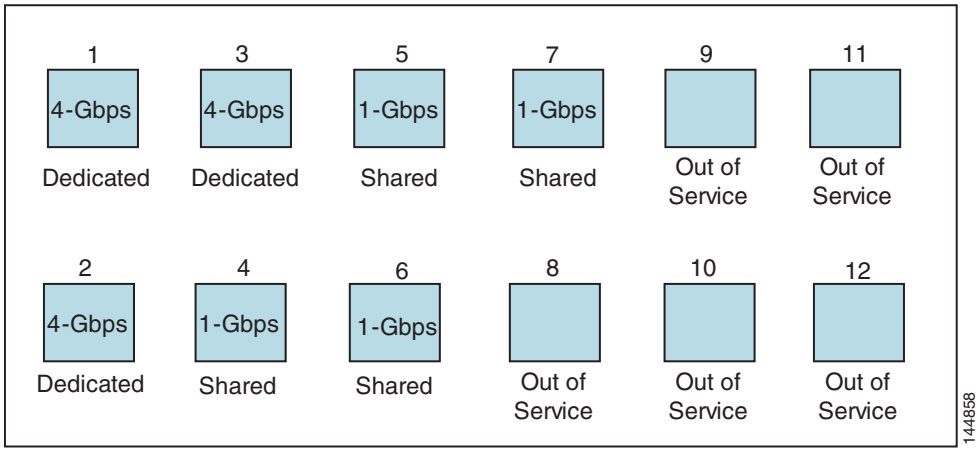

- Six ports with dedicated rate mode and 2-Gbps speed plus
  four ports with shared rate mode and 1-Gbps speed plus
  two ports put out-of-service (see Figure 14-4)

Note For an example of the configuration of the this example, see "Configuring a 48-port 4-Gbps Fibre Channel Switching Module Example" section on page 14-36.

Figure 14-4 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

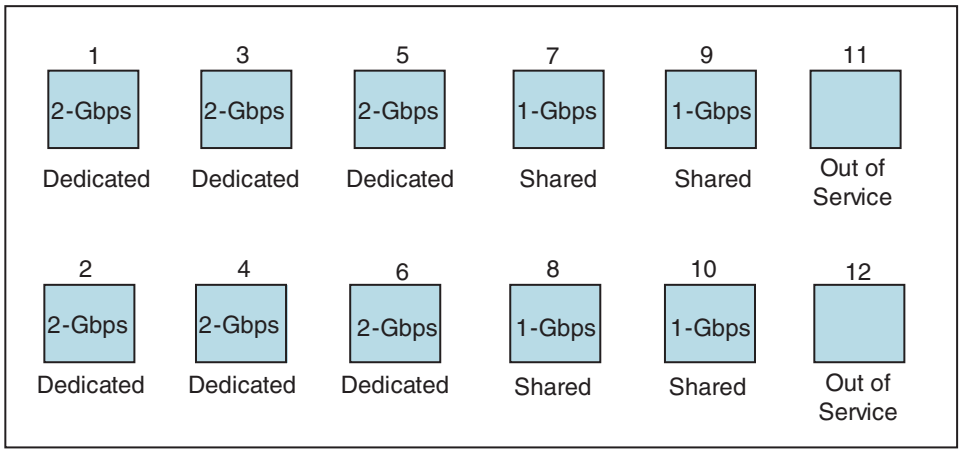

24-port 4-Gbps Fibre Channel Module BB_Credit Buffers

Table 14-6 lists the BB_credit buffer allocation for 24-port 4-Gbps Fibre Channel switching modules.

Table 14-6 24 Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults			
		Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
		ISL[1]	Fx Port	ISL[1]	Fx Port
User configurable BB_credit buffers	6000	250	16	16	16

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 24-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 4-Gbps Fibre Channel switching module consists of six ports. The ports in shared rate mode have bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth). The following example configurations are supported by the 24-port 4-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 4-Gbps speed (2:1 oversubscription) (default)
- Two ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 4-Gbps speed (with 4:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  three ports with dedicated rate mode and 2-Gbps speed plus
  two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)

- Six ports with dedicated rate mode and 2-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus three ports with shared rate mode and 1-Gbps speed (see Figure 14-5)

Note For an example of the configuration of the this example, see the "Configuring a 24-port 4-Gbps Fibre Channel Switching Module Example" section on page 14-36.

Figure 14-5 Example Speed and Rate Configuration on a 24-Port 4-Gbps Switching Module

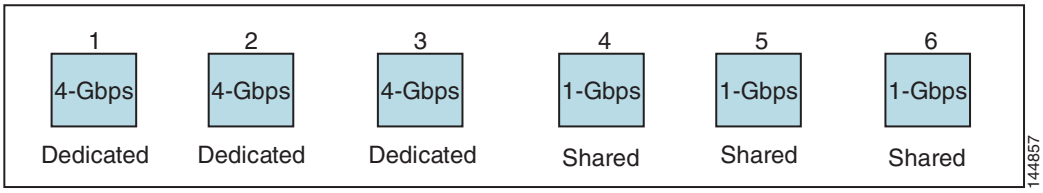

18-Port Fibre Channel/4-Port GigabitEthernet Multiservice Module BB_Credit Buffers

Table 14-6 lists the BB_credit buffer allocation for 18-port 4-Gbps multiservice modules.

Table 14-7 18-Port 4-Gbps Multiservice Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults			
		Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
		ISL[1]	Fx Port	ISL[1]	Fx Port
User configurable BB_credit buffers	4509	250	16	16	16

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on18-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

12-Port 4-Gbps Switching Module BB_Credit Buffers

Table 14-8 lists the BB_credit buffer allocation for 12-port 4-Gbps switching modules.

Table 14-8 12-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults	
		Dedicated Rate Mode 4-Gbps Speed	
		ISL [1]	Fx Port
User configurable BB_credit buffers	5488	250	16
Performance buffers	512 (shared)	145	12

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 12-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credit buffers after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 12)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

Note By default, the ports in the 12-port 4-Gbps switching modules come up in 4-Gbps dedicated rate mode but can be configured as 1-Gbps and 2-Gbps dedicated rate mode. Shared mode is not supported.

4-Port 10-Gbps Switching Module BB_Credit Buffers

Table 14-9 lists the BB_credit buffer allocation for 4-port 10-Gbps switching modules.

Table 14-9 4-port 10-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Module	BB_Credit Buffers Per Port Defaults	
		Dedicated Rate Mode 4-Gbps Speed	
		ISL[1]	F port[2]
User configurable BB_credit buffers	5488	250	16
Performance buffers	512 (shared)	145	12

1. ISL = E port or TE port.
2. Ports on the 4-port 10-Gbps cannot operate in FL port mode.

Note The ports in the 4-port 10-Gbps switching module only support 10-Gbps dedicated rate mode. FL port mode and shared rate mode are not supported.

The following considerations apply to BB_credit buffers on 4-port 10-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 4488 extra buffers available as extended BB_credits after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 4)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

BB_Credit Buffers for Fabric Switches

This section describes how buffer credits are allocated to Cisco MDS 9000 Fabric switches, and includes the following topics:

Cisco MDS 9134 Fabric Switch BB_Credit Buffers

Table 14-10 lists the BB_credit buffer allocation for 32-port 4-Gbps Fibre Channel switches.

Table 14-10 32-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL[1]	Fx Port
User configurable BB_credit buffers	64	64	64

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 32-port 4-Gbps switches:

- BB_credit buffers for connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers per port.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.

Cisco MDS 9124 Fabric Switch BB_Credit Buffers

Table 14-11 lists the BB_credit buffer allocation for 24-port 4-Gbps Fibre Channel switches.

Table 14-11 24-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL[1]	Fx Port
User configurable BB_credit buffers	64	16	16

1. ISL = E port or TE port.

Cisco MDS 9222i Multiservice Modular Switch BB_Credit Buffers

Table 14-12 lists the BB_credit buffer allocation for 18-port 4-Gbps Multiservice Modular switches.

Table 14-12 18-Port 4-Gbps Switching Module BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL[1]	Fx Port
User configurable BB_credit buffers	4509	250	16

1. ISL = E port or TE port.

Extended BB_Credits

Note Extended BB_credits are not supported on the Cisco MDS 9124 Fabric Switch, Cisco MDS 9134 Fabric Switch, Cisco MDS 9222i Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

To facilitate BB_credits for long haul links, the extended BB_credits feature allows the user to configure the receive buffers above the maximum value on all Generation 2 switching modules (see the "Buffer Credit Allocation" section on page 14-7). When necessary, you can reduce the buffers on one port and assign them to another port, exceeding the default maximum. The minimum extended BB_credits per port is 256 and the maximum is 4095.

In general, the user can configure any port in a port group to dedicated mode. To do this, you must first release the buffers from the other ports before configuring larger extended BB_credits for a port.

Note The ENTERPRISE_PKG license is required to use extended BB_credits on Generation 2 switching modules. Also, extended BB_credits are not supported by ports in shared rate mode.

All ports on the Generation 2 switching modules support extended BB_credits. There are no limitations

for how many extended BB_credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can take interfaces out of service to make more extended BB_credits available to other ports.

About Combining Generation 1 and Generation 2 Switching Modules

You can combine Generation 1 and Generation 2 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, are installed in the chassis.

Port Indexes

Cisco MDS 9000 switches allocate index identifiers for the ports on the modules. These port indexes cannot be configured. You can combine Generation 1 and Generation 2 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, are installed in the chassis.

Note On a switch with the maximum limit of 252 port index maximum limit, any new module that exceeds the limit when installed does not power up.

You can use the **show port index-allocation** command to display the allocation of port indexes on the switch.

```
switch# show port index-allocation

Module index distribution:
-------------------------------------------------------+
Slot | Allowed |      Alloted indices info             |
     |  range  | Total |          Index values         |
-----|---------|-------|-------------------------------|
 1   |   0- 255|   16  | 32-47                         |
 2   |   0- 255|   12  | 0-11                          |
```

```
 3   |   0- 255|   -   | (None)                    |
 4   |   0- 255|   -   | (None)                    |
 7   |   0- 255|   -   | (None)                    |
 8   |   0- 255|   -   | (None)                    |
 9   |   0- 255|   -   | (None)                    |
SUP  |  -----  |    3  | 253-255                   |
```

Generation 1 switching modules have specific numbering requirements. If these requirements are not met, the module does not power up. The port index numbering requirements include the following:

- If port indexes in the range of 256 to 1020 are assign to operational ports, Generation 1 switching modules do not power up.
- A block of contiguous port indexes is available. If such a block of port indexes is not available, Generation 1 modules do not power up. Table 14-13 shows the port index requirements for the Generation 1 modules.

Note If the switch has Supervisor-1 modules, the block of 32 contiguous port indexes must begin on the slot boundary. The slot boundary for slot 1 is 0, for slot 2 is 32, and so on. For Supervisor-2 modules, the contiguous block can start anywhere.

Table 14-13 Port Index Requirements for Generation 1 Modules

	Number of Port Indexes Required	
Generation 1 Module	**Supervisor-1 Module**	**Supervisor-2 Module**
16-port 2-Gbps Fibre Channel module	16	16
32-port 2-Gbps Fibre Channel module	32	32
8-port Gigabit Ethernet IP Storage Services module	32	32
4-port Gigabit Ethernet IP Storage Services module	32	16
32-port 2-Gbps Fibre Channel Storage Services Module (SSM).	32	32
14-port Fibre Channel/2-port Gigabit Ethernet Multiprotocol Services (MPS-14/2) module.	32	22

The allowed mix of Generation 1 and Generation 2 switching modules in a chassis is determined at run-time, either when booting up the switch or when installing the modules. In some cases, the sequence in which switching modules are inserted into the chassis determines if one or more modules is powered up. When a module does not power up because of a resource limitation, you can display the reason using **show module** command.

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    12     1/2/4 Gbps FC Module                                powered-dn
5    0      Supervisor/Fabric-2              DS-X9530-SF2-K9    active *

Mod  Power-Status  Power Down Reason
---  ------------  --------------------------
2    powered-dn    Insufficient resources (dest Index)
* this terminal session

Mod  MAC-Address(es)                          Serial-Num
```

```
---  ------------------------------------  ----------
1    00-0b-be-f7-4c-24 to 00-0b-be-f7-4c-28  JAB07030723
2    00-05-30-01-a8-b2 to 00-05-30-01-a8-b6  JAB090401AA
5    00-05-30-01-aa-7e to 00-05-30-01-aa-82  JAB091100TF

* this terminal session
```

The running configuration is updated when modules are installed. If you save the running configuration to the startup configuration (using the **copy running-config startup-config** command), during reboot the switch powers up the same set of modules as before the reboot regardless of the sequence in which the modules initialize. You can use the **show port index-allocation startup** command to display the index allocation the switch uses at startup.

```
switch# show port index-allocation startup

Startup module index distribution:
-----------------------------------------------------+
Slot | Allowed |      Alloted indices info            |
     |  range  | Total |          Index values        |
-----|---------|-------|------------------------------|
 1   |  -----  |   34  | 0-31,80-81                   |
 2   |  -----  |   32  | 32-63                        |
 3   |  -----  |   16  | 64-79                        |(Slot 1 shares 80-81)
 4   |  -----  |   48  | 96-127,224-239               |
SUP  | 253-255 |   3   | 253-255                      |
```

Note The output of the **show port index-allocation startup** command does not display anything in the "Allowed range" column because the command extracts the indices from the persistent storage service (PSS) and displaying an allowed range for startup indices is meaningless.

If a module fails to power up, you can use the **show module** *slot* **recovery-steps** command to display the reason. For information on recovering a module powered-down because port indexes are not available, refer to the *Cisco MDS 9000 Family Troubleshooting Guide, Release 3.x*.

Tip Whenever using mixed Generation 1 and Generation 2 modules, power up the Generation 1 modules first. During a reboot of the entire switch, the Generation 1 modules power up first (default behavior).

PortChannels

PortChannels have the following restrictions:

- The maximum number of PortChannels allowed is 256 if all switching modules are Generation 2.
- The maximum number of PortChannels allowed is 128 if the switching modules are Generation 1 or both Generation 1 and Generation 2.
- You must reserve the resources on the Generation 2 switching module interfaces to be used in the PortChannel.

Note The number of PortChannels allowed does not depend on the type of supervisor module.

When configuring PortChannels on switches with both Generation 1 and Generation 2 switching modules, configure the PortChannel and Generation 2 switching modules interfaces to auto with a maximum of 2 Gbps or configure the Generation 1 switching modules followed by the Generation 2 switching modules.

Note Generation 1 switching module interfaces do not support auto speed with max 2Gbps. Also, Generation 2 switching module interfaces cannot be forcefully added to a PortChannel if sufficient resources are not available.

Note Before adding a Generation 2 interface to a PortChannel, use the **show port-resources module** command to check for resource availability.

Table 14-14 describes the results of adding a member to a PortChannel for various configurations.

Table 14-14 PortChannel Configuration and Addition Results

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
No members	Any	Any	Generation 1 or Generation 2	Force	Pass
	Auto	Auto	Generation 1 or Generation 2	Normal or force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass or fail[1]
Generation 1 interfaces	Auto	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass or fail[1]
Generation 2 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass

1. Is resources not available.

Use the **show port-channel compatibility parameters** command to obtain information about PortChannel addition errors.

Configuring Generation 2 Module Interface Shared Resources

This section describes how to configure Generation 2 module interface shared resources and contains the following sections:

Displaying Interface Capabilities

Before configuring a Generation 2 interface, you can use the **show interface capabilities** command to display detailed information about the capabilities of the interface.

```
switch# show interface fc 9/1 capabilities
 Min Speed is 1 Gbps
 Max Speed is 4 Gbps
 FC-PH Version (high, low)                        (0,6)
 Receive data field size (max/min)                (2112/256)  bytes
 Transmit data field size (max/min)               (2112/128)  bytes
 Classes of Service supported are                 Class 2, Class 3, Class F
 Class 2 sequential delivery                      supported
 Class 3 sequential delivery                      supported
 Hold time (max/min)                              (100/1) micro sec
 BB state change notification                     supported
 Maximum BB state change notifications            14
 Rate Mode change                                 supported

 Rate Mode Capabilities                           Shared        Dedicated
  Receive BB Credit modification supported           yes             yes
  FX mode Receive BB Credit (min/max/default)     (1/16/16)      (1/250/16)
  ISL mode Receive BB Credit (min/max/default)        --         (2/250/250)
  Performace buffer modification supported            no              no

 Out of Service capable                            yes
 Beacon mode configurable                          yes
```

Configuration Guidelines for 48-Port and 24-Port 4-Gbps Fibre Channel Switching Modules

The 48-port and 24-port 4-Gbps Fibre Channel switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Shared and dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits

Migrating from Shared Mode to Dedicated Mode

To configure 48-port and 24-port 4-Gbps Fibre Channel switching modules when starting with the default configuration or when migrating from shared rate mode to dedicated rate mode, follow these guidelines:

1. Take unused interfaces out of service to release resources for other interfaces, if necessary.

   See the "Taking Interfaces Out of Service" section on page 14-33.

2. Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

   See the "Configuring Port Speed" section on page 14-23.

3. Configure the rate mode (dedicated or shared) to use.

   See the "Configuring Rate Mode" section on page 14-24.

4. Configure the port mode.

   See the "About Interface Modes" section on page 12-3.

   Note ISL ports cannot operate in shared rate mode.

5. Configure the BB_credits and extended BB_credits, as necessary.

   See the "About Buffer-to-Buffer Credits" section on page 12-33 and the "About Extended BB_credits" section on page 12-35.

Migrating from Dedicated Mode to Shared Mode

To configure 48-port and 24-port 4-Gbps Fibre Channel switching modules migrating from dedicated rate mode to shared rate mode, follow these guidelines:

1. Take unused interfaces out of service to release resources for other interfaces, if necessary.

   See the "Taking Interfaces Out of Service" section on page 14-33.

2. Configure the BB_credits and extended BB_credits, as necessary.

   See the "About Buffer-to-Buffer Credits" section on page 12-33 and the "About Extended BB_credits" section on page 12-35.

3. Configure the port mode.

   See the "About Interface Modes" section on page 12-3.

Note ISL ports cannot operate in shared rate mode.

4. Configure the rate mode (dedicated or shared) to use.

   See the "Configuring Rate Mode" section on page 14-24.

5. Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

   See the "Configuring Port Speed" section on page 14-23.

Configuration Guidelines for 12-Port 4-Gbps Switching Module Interfaces

The 12-port 4-Gbps switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these guidelines:

1. Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

   See the "Configuring Port Speed" section on page 14-23.

2. Configure the port mode.

   See the "About Interface Modes" section on page 12-3.

3. Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

   See the "About Buffer-to-Buffer Credits" section on page 12-33 and the "About Extended BB_credits" section on page 12-35.

Note If you change the port bandwidth reservation parameters on a 48-port or 24-port module, the change affects only the changed port. No other ports in the port group are affected.

Configuration Guidelines for 4-Port 10-Gbps Switching Module Interfaces

The 4-port 10-Gbps switching modules support the following features:

- Only 10-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and F port modes
- Extended BB_credits
- Performance buffers

Use the following guidelines to configure 4-port 10-Gbps switching modules when starting with the default configuration:

1. Configure the port mode.

   See the "About Interface Modes" section on page 12-3.

2. Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

   See the "About Buffer-to-Buffer Credits" section on page 12-33 and the "About Extended BB_credits" section on page 12-35.

Configuring Port Speed

The port speed on an interface, combined with the rate mode, determines the amount of shared resources available to the ports in the port group on a 48-port or 24-port 4-Gbps Fibre Channel switching module. Especially in the case of dedicated rate mode, the port group resources are reserved even though the bandwidth is not used. For example, if an interface is configured for autosensing (**auto**) and dedicated rate mode, then 4 Gbps of bandwidth is reserved even though the maximum operating speed is 2 Gbps. For the same interface, if autosensing with a maximum speed of 2 Gbps (**auto max 2000**) is configured, then only 2 Gbps of bandwidth is reserved and the unused 2 Gbps is shared with the other interface in the port group.

Caution Changing port speed and rate mode disrupts traffic on the port. Traffic on other ports in the port group is not affected.

Note The 4-port 10-Gbps switching module supports 10-Gbps traffic only.

To configure the port speed on an interface on a 4-Gbps switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport speed 4000`	Configures the port speed in megabits per second. Valid values are **1000**, **2000**, **4000** and **auto**. The **auto** parameter enables autosensing on the interface.
	`switch(config-if)# switchport speed auto`	Configures autosensing for the interface with 4 Gbps of bandwidth reserved.
	`switch(config-if)# switchport speed auto max 2000`	Configures autosensing with a maximum of 2 Gbps of bandwidth reserved.
	`switch(config-if)# no switchport speed`	Reverts to the default speed for the interface (**auto**).

Use the **show interface** command to verify the port speed configuration for an interface on a 24-port or 48-port 4-Gbps Fibre Channel switching module.

```
switch# show interface fc 9/1
```

```
fc9/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 22:01:00:05:30:01:9f:02
    Admin port mode is F
    snmp traps are enabled
    Port mode is F, FCID is 0xeb0002
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is shared
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      226 frames input, 18276 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      326 frames output, 21364 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 1 NOS, 0 loop inits
      3 output OLS, 2 LRR, 0 NOS, 0 loop inits
      16 receive B2B credit remaining
      64 transmit B2B credit remaining
```

Configuring Rate Mode

To configure the rate mode (dedicated or shared) on an interface on a 48-port or 24-port 4-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport rate-mode dedicated`	Reserves dedicated bandwidth for the interface. **Note** If you cannot reserve dedicated bandwidth on an interface, you might have exceeded the port group maximum bandwidth. Use the **show port-resources** command to determine what resources are already allocated.
	`switch(config-if)# switchport rate-mode shared`	Reserves shared (default) bandwidth for the interface.
	`switch(config-if)# no switchport rate-mode`	Reverts to the default state (**shared**).

Caution Changing port speed and rate mode disrupts traffic on the port.

Use **show port-resources module** command to verify the rate mode configuration for interfaces on a 48-port or 24-port 4-Gbps Fibre Channel switching module.

```
switch# show port-resources module 9
Module 9
Available dedicated buffers are 5400

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/1                                     16        4.0  shared
  fc9/2                                     16        4.0  shared
  fc9/3                                     16        4.0  shared
  fc9/4                                     16        4.0  shared
  fc9/5                                     16        4.0  shared
  fc9/6                                     16        4.0  shared

 Port-Group 2
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/7                                     16        4.0  shared
  fc9/8                                     16        4.0  shared
  fc9/9                                     16        4.0  shared
  fc9/10                                    16        4.0  shared
  fc9/11                                    16        4.0  shared
  fc9/12                                    16        4.0  shared

 Port-Group 3
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/13                                    16        4.0  shared
  fc9/14                                    16        4.0  shared
  fc9/15                                    16        4.0  shared
  fc9/16                                    16        4.0  shared
  fc9/17                                    16        4.0  shared
  fc9/18                                    16        4.0  shared

 Port-Group 4
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/19                                    16        4.0  shared
  fc9/20                                    16        4.0  shared
  fc9/21                                    16        4.0  shared
  fc9/22                                    16        4.0  shared
```

```
  fc9/23                                  16        4.0  shared
  fc9/24                                  16        4.0  shared
```

Configuring Oversubscription Ratio Restrictions

The 48-port and 24-port 4-Gbps Fibre Channel switching modules support oversubscription on switches with shared rate mode configurations. Table 14-15 describes the bandwidth allocation for oversubscribed interfaces configured in shared mode.

Table 14-15 Bandwidth Allocation for Oversubscribed Interfaces

Switching Module Type	Configured Speed	Reserved Bandwidth (Gbps)		Maximum Bandwidth (Gbps)
		Ratios enabled	Ratios disabled	
48 ports	Auto 4 Gbps	0.8	0.09	4
	Auto (maximum is 2 Gbps) 2 Gbps	0.4	0.045	2
	1 Gbps	0.2	0.0225	1
24 ports	Auto 4 Gbps	1	0.27	4
	Auto (maximum is 2 Gbps) 2 Gbps	0.5	0.135	2
	1 Gbps	0.25	0.067	1

By default, all 48-port and 24-port 4-Gbps Fibre Channel switching modules have restrictions on oversubscription ratios enabled.

As of Cisco SAN-OS Release 3.1(1) and later, you can disable restrictions on oversubscription ratios. All ports in 48-port and 24-port modules can be configured to operate at 4 Gbps in shared mode–even if other ports in the port group are configured in dedicated mode–regardless of available bandwidth. However, when oversubscription ratio restrictions are enabled you may not have all shared ports operating at 4 Gbps. For example, oversubscription ratios are enabled, and you have configured three 4 Gbps dedicated ports in one port group, no other ports in the same port group can be configured to operate at 4 Gbps.

```
switch# show port-resources module 8
Module 8
  Available dedicated buffers are 5478

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 0.8 Gbps
  Allocated dedicated bandwidth is 12.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc8/1                                     16        4.0  dedicated
  fc8/2                                     16        4.0  dedicated
  fc8/3                                     16        4.0  dedicated
  fc8/4 (out-of-service)
  fc8/5 (out-of-service)
  fc8/6 (out-of-service)
```

For dedicated ports, oversubscription ratio restrictions do not apply to the shared pool in port groups. So if oversubscription ratio restrictions are disabled, and you've configured three 4 Gbps dedicated ports in one port group, then you can configure all other ports in the same port group to operate at a shared rate of 4 Gbps. In the following example, a 24-port module has a group of 6 ports–3 dedicated ports are operating at 4 Gbps, and 3 shared ports operating at 4 Gbps:

```
switch# show port-resources module 8
Module 8
  Available dedicated buffers are 5382

Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 0.8 Gbps
  Allocated dedicated bandwidth is 12.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc8/1                                    16        4.0  dedicated
  fc8/2                                    16        4.0  dedicated
  fc8/3                                    16        4.0  dedicated
  fc8/4                                    16        4.0  shared
  fc8/5                                    16        4.0  shared
  fc8/6                                    16        4.0  shared

 Port-Group 2
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 0.8 Gbps
  Allocated dedicated bandwidth is 12.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc8/7                                    16        4.0  dedicated
  fc8/8                                    16        4.0  dedicated
  fc8/9                                    16        4.0  dedicated
  fc8/10                                   16        4.0  shared
  fc8/11                                   16        4.0  shared
  fc8/12                                   16        4.0  shared
...
```

When disabling restrictions on oversubscription ratios, all ports in shared mode on 48-port and 24-port 4-Gbps Fibre Channel switching modules must be shut down. When applying restrictions on oversubscription ratios, you must take shared ports out of service.

Note When restrictions on oversubscription ratios are disabled, the bandwidth allocation among the shared ports is proportionate to the configured speed. (If the configured speed is auto, then bandwidth is allocated assuming a speed of 4 Gbps.) For example, if you have three shared ports configured at 1, 2, and 4 Gbps, then the allocated bandwidth ratio is 1:2:4. In Cisco SAN-OS Release 3.0 and later (or when restrictions on oversubscription ratios are enabled), port bandwidths are allocated in equal proportions, regardless of port speed, so, the bandwidth allocation for the same three ports mentioned in the example would be 1:1:1.

Disabling Restrictions on Oversubscription Ratios

Before disabling restrictions on oversubscription ratios, ensure that you have explicitly shut down shared ports. To disable restrictions on oversubscription ratios on a 48-port or 24-port 4-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no rate-mode oversubscription-limit module 1`	Disables restrictions on oversubscription ratios for a module. **Note** You must enter this command separately for each module for which you want to remove the restrictions.
Step 3	`switch(config)# exit`	Exits configuration mode.
Step 4	`switch# copy running-config startup-config`	Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Use the **show running-config** command to view oversubscription ratios for a module. If oversubscription ratios are enabled, then no restriction appears in the output.

Example 14-1 Module with Restrictions on Oversubscription Ratios Disabled

```
switch# show running-config
version 3.1(1)
...
no rate-mode oversubscription-limit module 2
interface fc2/1
  switchport speed 2000
interface fc2/1
...
```

Oversubscription Ratio Restrictions Example

To disable restrictions on oversubscription ratios for ports on a 48-port Gen2 switch that is configured with both shared and dedicated ports, follow these steps:

Step 1 To disable restrictions on oversubscription ratios, you must shut down any shared ports. Use the **show port-resources** command to view the configuration on a module and to identify shared ports.

```
switch# show port-resources module 2
Module 2
Available dedicated buffers are 4656

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
```

```
  ----------------------------------------------------------------------
  fc2/1                                       16        4.0  shared
  fc2/2                                       16        4.0  shared
  fc2/3                                       16        4.0  dedicated
  fc2/4                                       16        4.0  shared
  fc2/5                                       16        4.0  shared
  fc2/6                                       16        4.0  dedicated
  fc2/7                                       16        4.0  dedicated
  fc2/8                                       16        4.0  shared
  fc2/9                                       16        4.0  shared
  fc2/10                                      16        4.0  shared
  fc2/11                                      16        4.0  shared
  fc2/12                                      16        4.0  shared
  ...

 Port-Group 4
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ----------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ----------------------------------------------------------------------
  fc2/37                                      16        4.0  shared
  fc2/38                                      16        4.0  shared
  fc2/39                                      16        4.0  dedicated
  fc2/40                                      16        4.0  dedicated
  fc2/41                                      16        4.0  dedicated
  fc2/42                                      16        4.0  shared
  fc2/43                                      16        4.0  shared
  fc2/44                                      16        4.0  shared
  fc2/45                                      16        4.0  shared
  fc2/46                                      16        4.0  shared
  fc2/47                                      16        4.0  shared
  fc2/48                                      16        4.0  shared
```

Step 2 Shut down all shared ports for which you want to remove restrictions on oversubscription ratios.

```
switch (config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch (config-if)# shutdown
```

Step 3 Display the interface status to confirm the shutdown of all shared ports.

```
switch(config-if)# end
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper    Port
                  Mode   Trunk                         Mode  Speed   Channel
                         Mode                                 (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      down            swl    --          --
fc2/2      1      FX     --      down            swl    --          --
fc2/3      1      T      --      up              swl    --          --
fc2/4      1      FX     --      down            swl    --          --
fc2/5      1      FX     --      down            swl    --          --
fc2/6      1      TE     --      up              swl    --          --
fc2/7      1      TE     --      up              swl    --          --
fc2/8      1      FX     --      down            swl    --          --
...
fc2/48     1      FX     --      down            swl    --          --
```

Step 4 Disable restrictions on oversubscription ratios for the ports.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
```

```
switch(config)# no rate-mode oversubscription-limit module 2
```

Step 5 Bring up the ports that you shut down in step 2, and display their status to confirm that they are no longer shut down.

```
switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch(config-if)# no shutdown
switch(config-if)# end
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                                (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      up              swl    --           --
fc2/2      1      FX     --      up              swl    --           --
fc2/3      1      T      --      up              swl    --           --
fc2/4      1      FX     --      up              swl    --           --
fc2/5      1      FX     --      up              swl    --           --
fc2/6      1      TE     --      up              swl    --           --
fc2/7      1      TE     --      up              swl    --           --
fc2/8      1      FX     --      up              swl    --           --
...
fc2/48     1      FX     --      up              swl    --           --
```

Step 6 Confirm that the ports are now operating with no restrictions on oversubscription ratios.

```
switch# show running-config | include oversubscription-limit
no rate-mode oversubscription-limit module 2 <---indicates no restrictions on
                                              oversubscrption ratios
```

Step 7 Save the new oversubscription ratio configuration to the startup configuration.

```
switch# copy running-config startup-config
```

Enabling Restrictions on Oversubscription Ratios

Caution You must enable restrictions on oversubscription ratios before you can downgrade modules to a previous release.

Before enabling restrictions on oversubscription ratios, ensure that you have explicitly configured shared ports to out-of-service mode. To enable restrictions on oversubscription ratios on a 48-port or 24-port 4-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48`	Specifies the port interfaces for which you want to enable restrictions on oversubscription ratios.
Step 3	`switch(config-if)# shutdown`	Shuts down shared ports.
Step 4	`switch(config-if)# out-of-service`	Takes shared ports out of service.

	Command	Purpose
Step 5	`switch# rate-mode oversubscription-limit module 1`	Enables restrictions on oversubscription ratios for the module. **Note** You must enter this command separately for each module for which you want to add the restriction.

Step 6	`switch# config t`
	`switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48`
	`switch(config-if)# no out-of-service`
	`switch(config-if)# no shutdown`
		Returns all shared ports to service.

Step 7	`switch(config)# exit`	Exits configuration mode.
Step 8	`switch# copy running-config startup-config`	Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Configuring Bandwidth Fairness

As of Cisco SAN-OS Release 3.1(2) and later, all 48-port and 24-port 4-Gbps Fibre Channel switching modules, as well as 18-port Fibre Channel/4-port GigabitEthernet Multiservice modules, have bandwidth fairness enabled by default, which improves fairness of bandwidth allocation among all ports and provides better throughput average to individual data streams. Bandwidth fairness can be configured per module.

Caution When you disable or enable bandwidth fairness, the change does not take effect until you reload the module.

Use the **show module bandwidth-fairness** command to check whether ports in a module are operating with bandwidth fairness enabled or disabled.

```
switch# show module 2 bandwidth-fairness
Module 2 bandwidth-fairness is enabled
```

Note This feature is only supported onthe 48-port and 24-port 4-Gbps Fibre Channel switching modules, as well as the 18-port Fibre Channel/4-port GigabitEthernet Multiservice module.

Enabling Bandwidth Fairness

To enable bandwidth fairness on switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rate-mode bandwidth-fairness module 1`	Enables bandwidth fairness for a module. **Note** You must enter this command separately for each module for which you want to enable bandwidth fairness. You must reload the module for the command to take effect.
Step 3	`switch(config)# exit`	Exits configuration mode.

Disabling Bandwidth Fairness

Note If you disable bandwidth fairness, up to a 20 percent increase in internal bandwidth allocation is possible for each port group; however, bandwidth fairness is not guaranteed when there is a mix of shared and full-rate ports in the same port group.

To disable bandwidth fairness on a switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no rate-mode bandwidth-fairness module 1`	Disables bandwidth fairness for a module. **Note** You must enter this command separately for each module for which you want to disable bandwidth fairness. You must reload the module for the command to take effect.
Step 3	`switch(config)# exit`	Exits configuration mode.

Upgrade or Downgrade Scenario

When you are upgrading from a release earlier than Cisco SAN-OS Release 3.1(2), all modules operate with bandwidth fairness disabled until the next module reload. After the upgrade, any new module that is inserted has bandwidth fairness enabled.

When you are downgrading to a release earlier than Cisco SAN-OS Release 3.1(2), all modules keep operating in the same bandwidth fairness configuration prior to the downgrade. After the downgrade, any new module that is inserted has bandwidth fairness disabled.

Note After the downgrade, any insertion of a module or module reload will have bandwidth fairness disabled.

Taking Interfaces Out of Service

You can take interfaces out of service on Generation 2 switching modules. When an interface is out of service, all the shared resources for the interface are released as well as the configuration associated with those resources.

Note The interface must be disabled using a **shutdown** command before it can be taken out of service.

Caution Taking interfaces out of service releases all the shared resources to ensure that they are available to other interfaces. This causes the configuration in the shared resources to revert to default when the interface is brought back into service. Also, an interface cannot come back into service unless the default shared resources for the port are available. The operation to free up shared resources from another port is disruptive.

Note The interface cannot be a member of a PortChannel.

To take an interface out of service, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# no channel-group`	Removes the interface from a PortChannel.
Step 4	`switch(config-if)# shutdown`	Disables the interface.

Step 5	`switch(config-if)# out-of-service`
	`Putting an interface into out-of-service will cause its shared resource configuration to revert to default`
	`Do you wish to continue(y/n)? [n] y`
		Takes the interface out of service.

Use the **show port-resources module** command to verify the out-of-service configuration for interfaces on a Generation 2 switching module.

```
switch# show port-resources module 9
Module 9
Available dedicated buffers are 5429

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/1                                     16        4.0  shared
  fc9/2 (out-of-service)
  fc9/3                                     16        4.0  shared
  fc9/4                                     16        4.0  shared
  fc9/5                                     16        4.0  shared
```

```
  fc9/6                                    16        4.0  shared
...
```

Releasing Shared Resources in a Port Group

When you want to reconfigure the interfaces in a port group on a Generation 2 module, you can return the port group to the default configuration to avoid problems with allocating shared resources.

Note The interface cannot be a member of a PortChannel.

Caution Releasing shared resources disrupts traffic on the port. Traffic on other ports in the port group is not affected.

To release the shared resources for a port group, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode. **Tip** You can use an interface range to release the resources for all interfaces in a port group.
Step 3	`switch(config-if)# no channel-group`	Removes the interface from a PortChannel.
Step 4	`switch(config-if)# shutdown`	Disables the interface.

Step 5	`switch(config-if)# out-of-service`
	`Putting an interface into out-of-service will cause its shared resource configuration to revert to default`
	`Do you wish to continue(y/n)? [n] y`
		Takes the interface out of service.

Step 6	`switch(config-if)# no out-of-service`	Makes the interface available for service. Repeat Step 2 through Step 6 for all the interfaces in the port group.

Enabling the Buffer-to-Buffer State Change Number

The BB_SC_N field (word 1, bits 15-12) specifies the buffer-to-buffer state change (BB_SC) number. The BB_SC_N field indicates that the sender of the port login (PLOGI) or fabric login (FLOGI) frame is requesting twice the number of frames specified by BB_SC_N to be sent between two consecutive BB_SC send primitives, and twice the number of R_RDY primitives to be sent between two consecutive BB_SC receive primitives.

To use the BB_SC_N field during PLOGI or FLOGI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcbbscn`	Enables the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.
	`switch(config-if)# no switchport fcbbscn`	Disables (default) the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.

Disabling ACL Adjacency Sharing for System Image Downgrade

As of Cisco MDS SAN-OS Release 3.0(3), Fibre Channel ACL adjacency sharing is enabled by default on the switches with an active Generation 2 switching module. Fibre Channel ACL adjacency sharing improves the performance for zoning and inter-VSAN routing (IVR) network address translation (NAT). To prevent disruptions when downgrading the system image on your switch to a release prior to Cisco SAN-OS Release 3.0(3), issue the following command in EXEC mode:

```
switch# system no acl-adjacency-sharing
```

To reenable Fibre Channel ACL adjacency sharing on your switch, issue the following command in EXEC mode:

```
switch# system acl-adjacency-sharing
```

Displaying SFP Diagnostic Information

You can use the **show interface transceiver** command to display small form-factor pluggable (SFP) diagnostic information for Generation 2 switching modules.

```
switch# show interface transceiver
...
fc12/12 sfp is present
    name is CISCO-FINISAR
    part number is FTRJ-8519-7D2CS1
    revision is A
    serial number is H11TVQB
    fc-transmitter type is short wave laser w/o OFC (SN)
    fc-transmitter supports intermediate distance link length
    media type is multi-mode, 62.5m (M6)
    Supported speed is 200 MBytes/sec
    Nominal bit rate is 2100 MBits/sec
    Link length supported for 50/125mm fiber is 300 m(s)
    Link length supported for 62.5/125mm fiber is 150 m(s)
    cisco extended id is unknown (0x0)

    no tx fault, rx loss, no sync exists, Diag mon type 104
    SFP Diagnostics Information
        Temperature      :  24.33 Celsius
        Voltage          :   3.33 Volt
        Current          :   0.04 mA       --
```

```
        Optical Tx Power :  N/A   dBm      --
        Optical Rx Power :  N/A   dBm      -
    Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
...
```

Example Configurations

This section describes example configurations and includes the following sections:

Configuring a 24-port 4-Gbps Fibre Channel Switching Module Example

This section describes how to configure the example shown in Figure 14-5 on page 14-12.

Step 1 Select interfaces fc 3/1 through fc 3/3.

```
switch# config t
switch(config)# interface fc 3/1 - 3
```

Step 2 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode e
```

Step 3 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 3/4 through fc 3/6.

```
switch# config t
switch(config)# interface fc 3/4 - 6
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 1000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuring a 48-port 4-Gbps Fibre Channel Switching Module Example

This section describes how to configure the example shown in Figure 14-4 on page 14-11.

Step 1 Select interfaces fc 4/11 through fc 4/12.

```
switch# config t
switch(config)# interface fc 4/11 - 12
```

Step 2 Disable the interfaces and take them out of service.

```
switch(config-if)# shutdown
switch(config-if)# out-of-service
```

Step 3 Return to configuration mode.

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

Step 4 Select the interfaces fc 4/1 through fc 4/6.

```
switch# config t
switch(config)# interface fc 4/1 - 6
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed auto max 2000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode e
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 7 Select the interfaces fc 4/7 through fc 4/10.

```
switch# config t
switch(config)# interface fc 4/7 - 10
```

Step 8 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 1000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 9 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Default Settings

Table 14-16 lists the default settings for Generation 2 interface parameters.

Table 14-16 Default Generation 2 Interface Parameters

Parameter	Default			
	48-Port 4-Gbps Switching Module	**24-Port 4-Gbps Switching Module**	**12-Port 4-Gbps Switching Module**	**4-Port 10-Gbps Switching Module**
Speed mode	**auto**[1]	**auto**[1]	**auto**[1]	**auto**[2]
Rate mode	**shared**	**shared**	**dedicated**	**dedicated**
Port mode	Fx	Fx	auto[3]	auto[4]
BB_credit buffers	16	16	250	250
Performance buffers	–	–	145[5]	145[5]

1.
2. The 4-port 10-Gbps switching module only supports 10-Gbps traffic.
3. Auto port mode on the 12-port 4-Gbps switching module interfaces can operate in E port mode, TE port mode, and Fx port mode.
4. Auto port mode on the 4-port 10-Gbps switching module interfaces can operate in E port mode, TE port mode, and F port mode.
5. Performance buffers are shared among all ports on the module.

CHAPTER 15

Configuring Trunking

This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format (see Figure 15-1).

Figure 15-1 Trunking

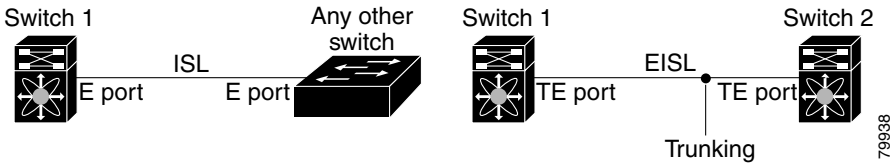

The trunking feature includes the following restrictions:

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

Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c_Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Trunking Configuration Guidelines

If you misconfigure VSAN configurations across E ports, you could face consequences such as merging the traffic in two VSANs (thus causing both VSANs to mismatch). The trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid merging VSANs (see Figure 15-2).

Figure 15-2 VSAN Mismatch

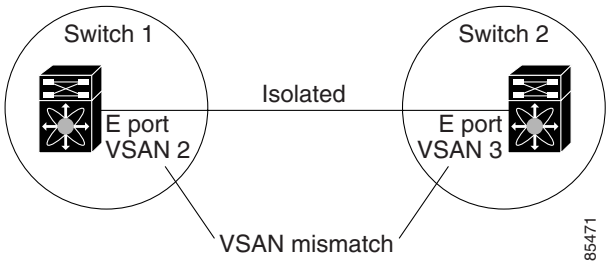

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see Figure 15-3).

Figure 15-3 Third-Party Switch VSAN Mismatch

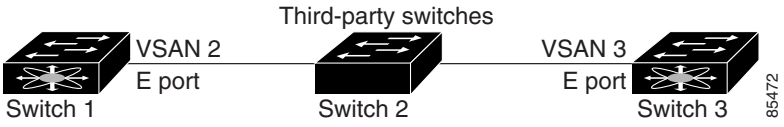

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Trunking Protocol

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

- 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 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). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

Note We recommend that both ends of a trunking ISL belong to the same port VSAN. On certain platforms or fabric switches where the port VSANs are different, one end returns an error, and the other is not connected.

Tip To avoid inconsistent configurations, disable all E ports with a **shutdown** command before enabling or disabling the trunking protocol.

This section explains how to configure trunking and contains the following topics:

Enabling or Disabling the Trunking Protocol

To enable or disable the trunking protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# no trunk protocol enable` `switch(config)#`	Disables the trunking protocol.
	`switch(config)# trunk protocol enable` `switch(config)#`	Enables trunking protocol (default).

About Trunk Mode

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configuration at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends (see Table 15-1).

Table 15-1 Trunk Mode Status Between Switches

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

Tip The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other set to on.

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

Configuring Trunk Mode

To configure trunk mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# switchport trunk mode on`	Enables (default) the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode off`	Disables the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode auto`	Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

About Trunk-Allowed VSAN Lists

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the 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 Figure 15-4, 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 Figure 15-4.

Figure 15-4 Default Allowed-Active VSAN Configuration

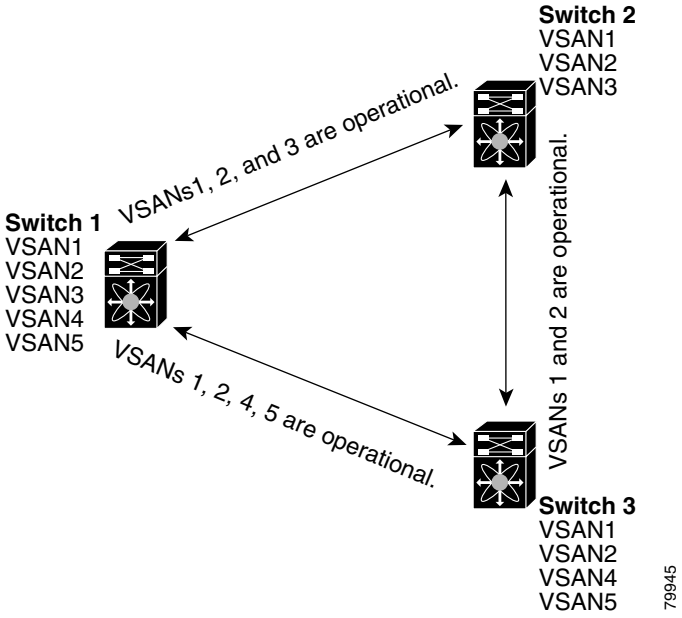

You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using Figure 15-4 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 15-5). 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 shall include VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 shall include VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 shall include VSAN 1, 2, and 5.

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

Figure 15-5 Operational and Allowed VSAN Configuration

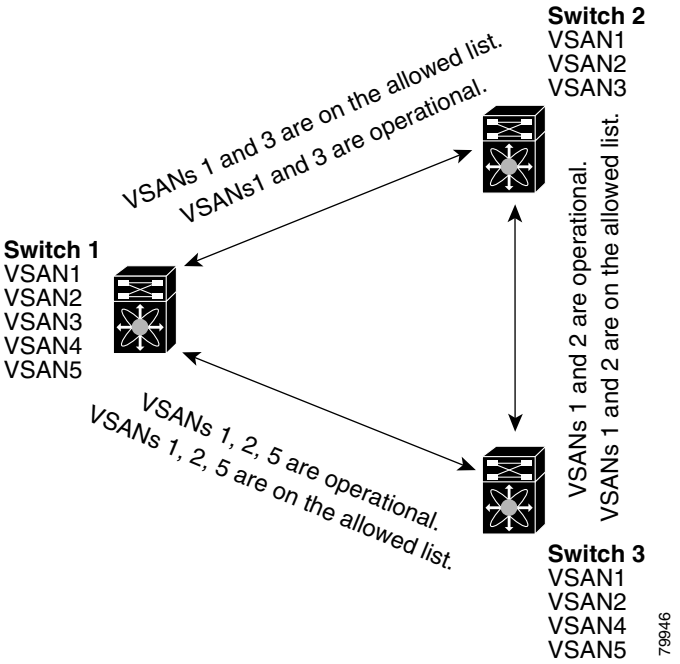

Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# switchport trunk allowed vsan 2-4`	Changes the allowed list for the specified VSANs.
	`switch(config-if)# switchport trunk allowed vsan add 5` `updated trunking membership`	Expands the specified VSAN (5) to the new allowed list.
	`switch(config-if)# no switchport trunk allowed vsan 2-4`	Deletes VSANs 2, 3, and 4.
	`switch(config-if)# no switchport trunk allowed vsan add 5`	Deletes the expanded allowed list.

Displaying Trunking Information

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples 15-1 to 15-3.

Example 15-1 Displays a Trunked Fibre Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
```

```
    Hardware is Fibre Channel
    Port WWN is 20:0d:00:05:30:00:58:1e
    Peer port WWN is 20:0d:00:05:30:00:59:1e
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Receive B2B Credit is 255
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      233996 frames input, 14154208 bytes, 0 discards
        0 CRC, 0 unknown class
        0 too long, 0 too short
      236 frames output, 13818044 bytes, 0 discards
      11 input OLS, 12 LRR, 10 NOS, 28 loop inits
      34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

Example 15-2 Displays the Trunking Protocol

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

Example 15-3 Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
    Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
    Vsan 1 is up, FCID is 0x760001
    Vsan 2 is up, FCID is 0x6f0001

fc3/11 is trunking
    Belongs to port-channel 6
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
```

Default Settings

Table 15-2 lists the default settings for trunking parameters.

Table 15-2 Default Trunk Configuration Parameters

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

CHAPTER 16

Configuring PortChannels

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

This chapter discusses the PortChannel feature provided in the switch and includes the following sections:

About PortChannels

A PortChannel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

Note See the "Fail-Over Scenarios for PortChannels and FSPF Links" section on page 25-3 for fail-over scenarios.

Cisco MDS 9000 Family switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, support a maximum of 128 PortChannels. Switches with only a Generation 2 switching module support a maximum of 256 PortChannels (with 16 interfaces per PortChannel). A PortChannel number refers to the unique (to each switch) identifier associated with each channel group. This number ranges from of 1 to 256.

This section describes PortChannels and contains the following topics:

PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. Figure 16-1 illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

Figure 16-1 PortChannel Flexibility

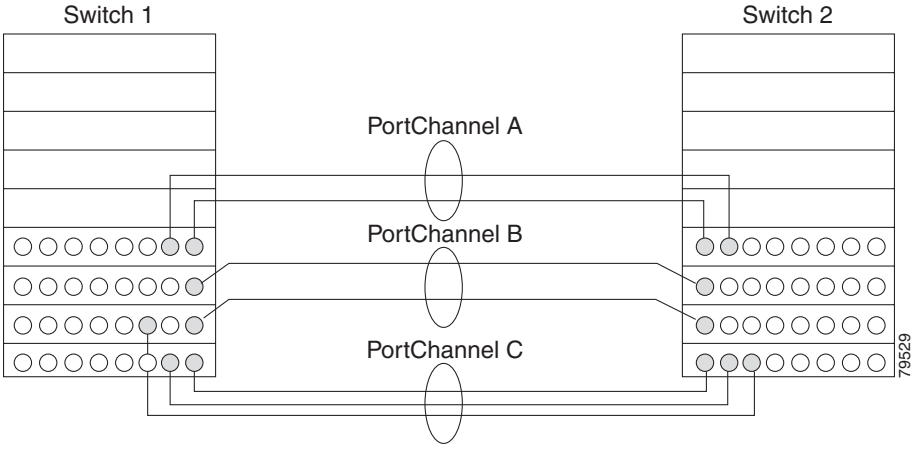

32-Port Switching Module Configuration Guidelines

The 32-port switching module guidelines apply to the following hardware:

- 32-port 2-Gbps or 1-Gbps switching modules
- Cisco MDS 9140 switches

When configuring these host-optimized ports, the following PortChannel guidelines apply:

- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules—only the first port of each group of 4 ports is included in a PortChannel.
  - You can configure only the first port in each 4-port group (for example, the first port in ports 1–4, the fifth port in ports 5–8, and so on) as an E port. If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2–4, 6–8, and so on) are not usable and remain in the shutdown state.
  - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.

Note In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the Cisco MDS 9120 Switch and 8 ports in the Cisco MDS 9140 Switch) are full line rate like the 16-port switching module. The other ports (16 ports in the Cisco MDS 9120 Switch and 32 ports in the Cisco MDS 9140 Switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

About PortChanneling and Trunking

Trunking is a commonly used storage industry term. However, the Cisco SAN-OS software and switches in the Cisco MDS 9000 Family implement trunking and PortChanneling as follows:

- PortChanneling enables several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. When trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see Figure 16-2 and Figure 16-3).

See Chapter 15, "Configuring Trunking," for information on trunked interfaces.

Figure 16-2 Trunking Only

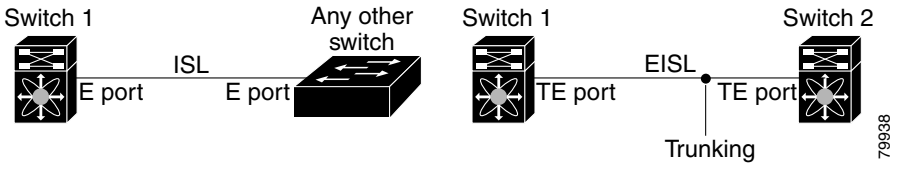

PortChanneling and trunking are used separately across an ISL:

Figure 16-3 PortChanneling and Trunking

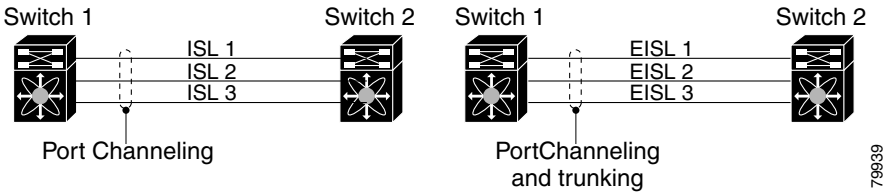

- PortChanneling—Interfaces can be channeled between E ports and TE ports.
- Trunking—Interfaces can be trunked only between TE ports. Trunking permits carrying traffic on multiple VSANs between switches.

  See Chapter 19, "Configuring and Managing VSANs."

Both PortChanneling and trunking can be used between TE ports over EISLs.

About Load Balancing

Two mechanisms support the load balancing functionality:

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

Figure 16-4 illustrates how source ID 1 (SID1) and destination ID1 (DID1) based load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 16-4 SID1 and DID1 Based Load Balancing

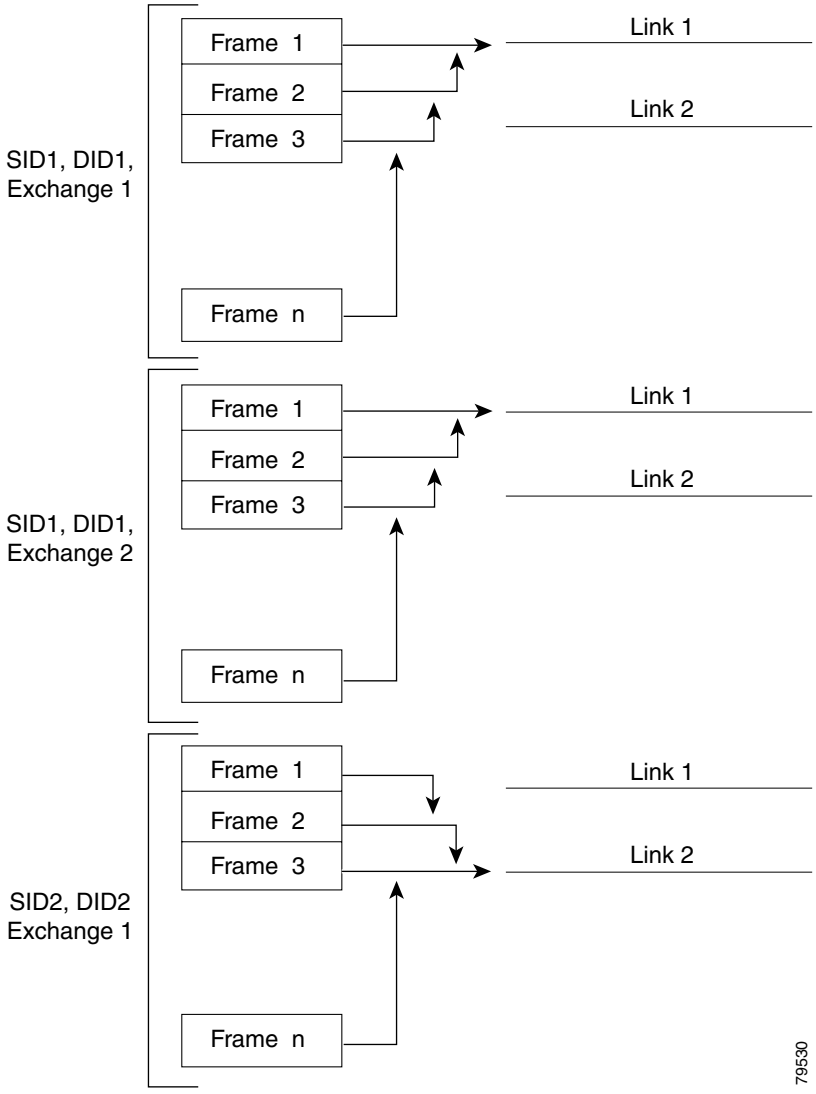

Figure 16-5 illustrates how exchange based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 16-5 SID1, DID1, and Exchange Based Load Balancing

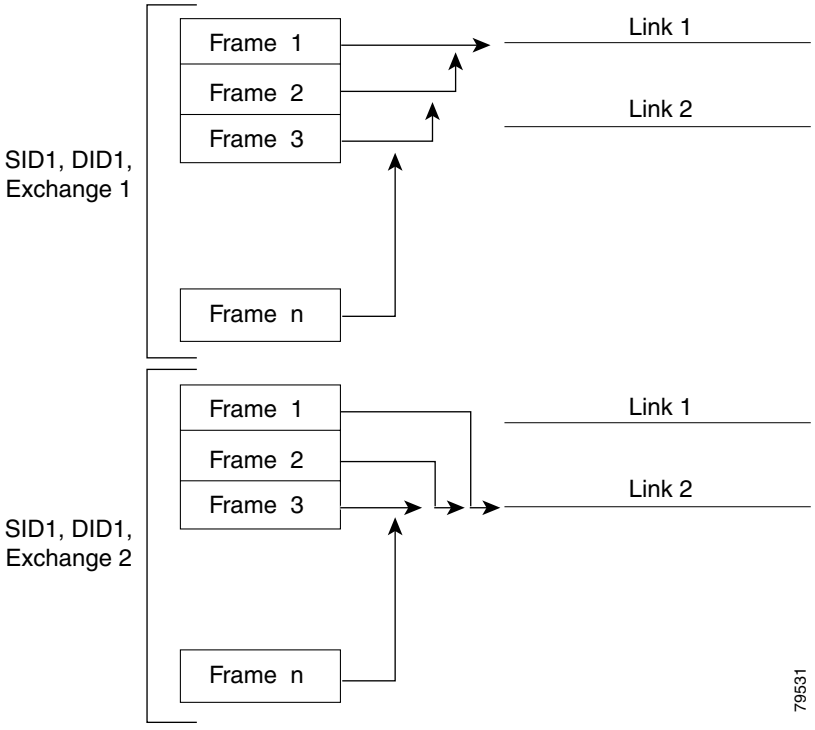

For more information on configuring load balancing and in-order delivery features, see the "Operational State of a VSAN" section on page 19-9.

PortChannel Configuration

PortChannels are created with default values. You can change the default configuration just like any other physical interface.

Figure 16-6 provides examples of valid PortChannel configurations.

Figure 16-6 Valid PortChannel Configurations

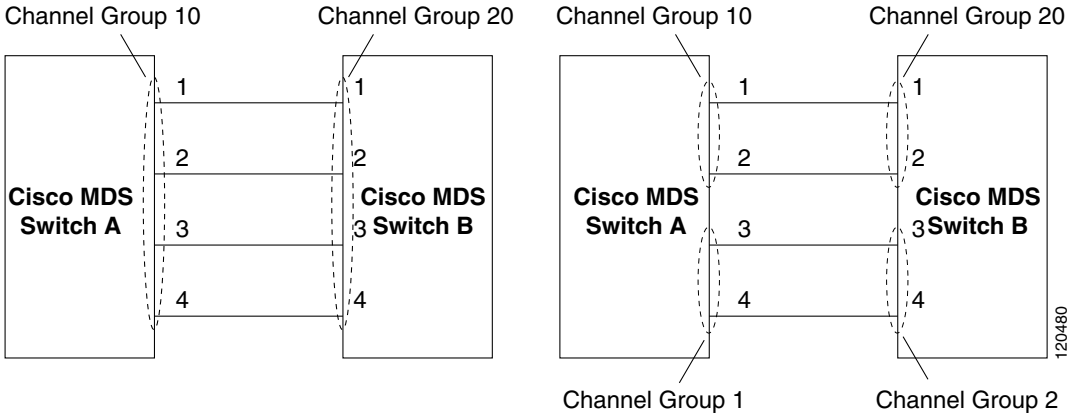

Figure 16-7 provides examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 16-7 Misconfigured Configurations

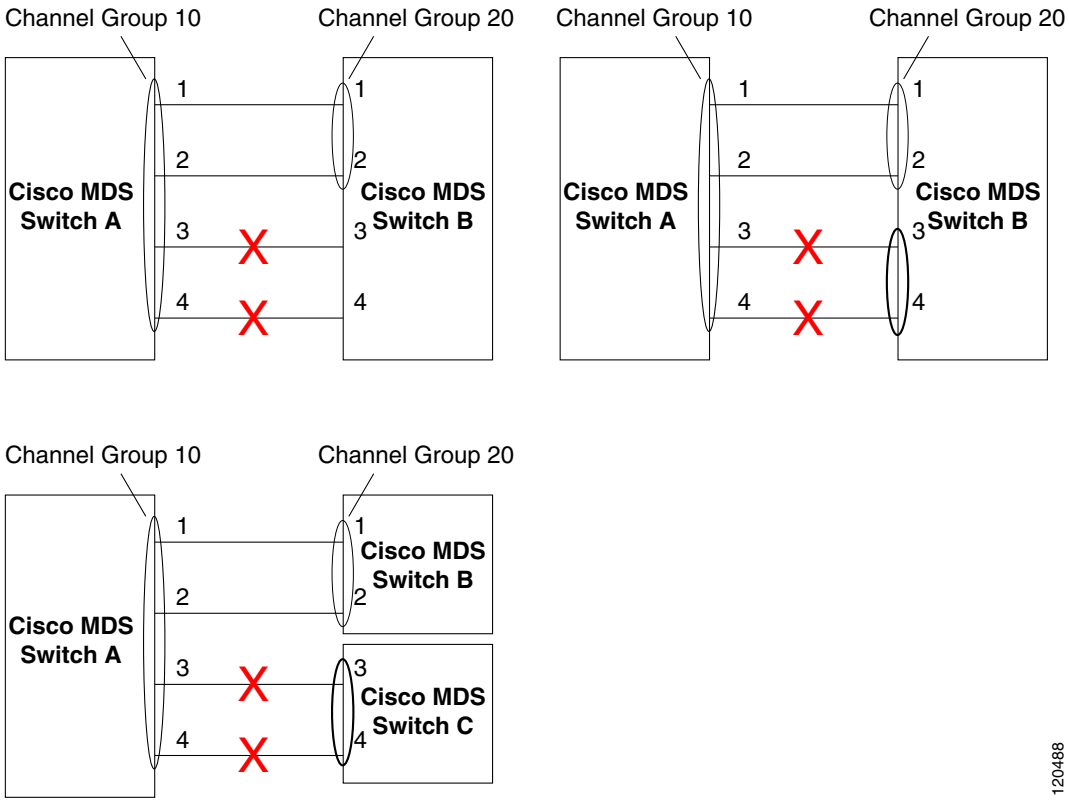

This section shows how to configure and modify PortChannels and contains the following topics:

About PortChannel Configuration

Before configuring a PortChannel, consider the following guidelines:

- Configure the PortChannel across switching modules to prevent redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.

Note On switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, you can configure a maximum of 128 PortChannels. On switches with only Generation 2 switching modules, you can configure a maximum of 256 PortChannels.

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled because an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see Figure 16-7 for an example of an invalid configuration).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and reenable the links.

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

Issue the **show interface** command for that interface to verify that the PortChannel is functioning as required.

Creating a PortChannel

To create a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface port-channel 1` `switch(config-if)#`	Configures the specified PortChannel (1) using the default ON mode.

About PortChannel Modes

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

- ON (default)—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON mode require you to explicitly enable and disable the PortChannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- ACTIVE—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. The ACTIVE PortChannel mode allows automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 16-1 compares ON and ACTIVE modes.

Table 16-1 Channel Group Configuration Differences

ON Mode	ACTIVE Mode
No protocol is exchanged.	A PortChannel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel.	Moves interfaces to the isolated state if its operational values are incompatible with the PortChannel.
When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end.	When you add or modify a PortChannel interface, the PortChannel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a PortChannel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

To configure active mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface port-channel 1` `switch(config-if)#`	Configures the specified PortChannel (1) using the default ON mode.
Step 3	`switch(config-if)# channel mode active`	Configures the ACTIVE mode.
	`switch(config-if)# no channel mode active`	Reverts to the default ON mode.

About PortChannel Deletion

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Graceful Shutdown" section on page 12-12).

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

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

Deleting PortChannels

To delete a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# no interface port-channel 1`
	`port-channel 1 deleted and all its members disabled`
	`please do the same operation on the switch at the other end of the port-channel`
	`switch(config)#`
		Deletes the specified PortChannel (1), its associated interface mappings, and the hardware associations for this PortChannel.

Interfaces in a PortChannel

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

This section describes interface configuration for a PortChannel and includes the following topics:

Note For information about PortChannel support on Generation 2 switching modules, see the "PortChannels" section on page 14-18.

About Interface Addition to a PortChannel

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

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "32-Port Switching Module Configuration Guidelines" section on page 16-2 and "Graceful Shutdown" section on page 12-12).

Compatibility Check

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

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

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

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

Suspended and Isolated States

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

- An interface enters the suspended state if the interface is configured in the ON mode.
- An interface enters the isolated state if the interface is configured in the ACTIVE mode.

See the "Reason Codes" section on page 12-8.

Adding an Interface to a PortChannel

To add an interface to a PortChannel, follow these steps

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/15` `switch(config-if)#`	Configures the specified port interface (fc1/15).
Step 3	`switch(config-if)# channel-group 15` `fc1/15 added to port-channel 15 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up`	Adds physical Fibre Channel port 1/15 to channel group 15. If channel group 15 does not exist, it is created. The port is shut down.

To add a range of ports to a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# interface fc1/1 - 5` `switch(config-if)#`	Configures the specified range of interfaces. In this example, interfaces from 1/1 to 1/5 are configured.

Step 3	`switch(config-if)# channel-group 2`
	`fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 added to port-channel 2 and disabled`
	`please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up`
		Adds physical interfaces 1/1, 1/2, 1/3, 1/4, and 1/5 to channel group 2. If channel group 2 does not exist, it is created.
		If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces.

Forcing an Interface Addition

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

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

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

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Graceful Shutdown" section on page 12-12).

To force the addition of a port to a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Specifies the interface fc1/1.

Step 3	`switch(config-if)# channel-group 1 force`
	`fc1/1 added to port-channel 1 and disabled`
	`please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up`
	`switch(config-if)#`
		Forces the addition of the physical port for interface fc1/1 to channel group 1. The E port is shut down.

About Interface Deletion from a PortChannel

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

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

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "32-Port Switching Module Configuration Guidelines" section on page 16-2 and "Graceful Shutdown" section on page 12-12).

Deleting an Interface from a PortChannel

To delete a physical interface (or a range of physical interfaces) from a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch(config)# interface fc1/1` `switch(config-if)#`	Enters the selected physical interface level.
	`switch(config)# interface fc1/1 - 5` `switch(config-if)#`	Enters the selected range of physical interfaces.

Step 2	`switch(config-if)# no channel-group 2`
	`fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 removed from port-channel 2 and disabled. Please do the same operation on the switch at the other end of the port-channel`
	`switch(config-if)#`
		Deletes the physical Fibre Channel interfaces in channel group 2.

PortChannel Protocol

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco SAN-OS software provides robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

A protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default.

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

The PortChannel protocol uses two sub-protocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications like write acceleration work for PortChannels over FCIP links.
- Autocreation protocol—Automatically aggregates compatible ports into a PortChannel.

This section describes how to configure the PortChannel protocol and includes the following sections:

About Channel Group Creation

Note Channel groups are not supported on internal ports in the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

Assuming link A1-B1 comes up first in Figure 16-8, that link is operational as an individual link.When the next link, say A2-B2 comes up, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (and hence, the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is because of the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 16-8 Autocreating Channel Groups

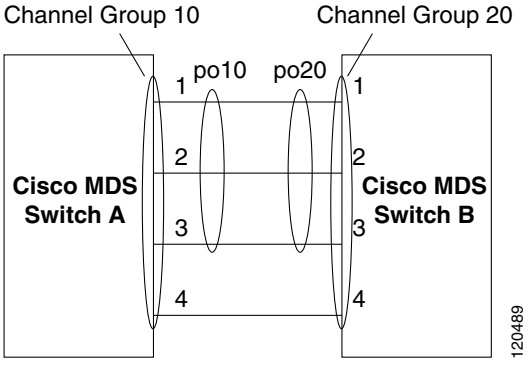

The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

Table 16-2 identifies the differences between user-configured and auto-configured channel groups.

Table 16-2 Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.
You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration.	All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.
Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface.	Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist.

About Autocreation

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated PortChannel.
  - A port is aggregated with another compatible port to form a new PortChannel.
- Newly created PortChannels are allocated from the maximum possible PortChannel (128 for Generation 1 or a combination of Generation 1 and Generation 2 switches, or 256 for Generation 2 switches) in a decreasing order based on availability. If all 128 (or 256) numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel.
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.

- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.
- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.

Tip When enabling autocreation in any switch in the Cisco MDS 9000 Family, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to an autocreated PortChannel.

Enabling and Configuring Autocreation

To configure automatic channel groups, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/13` `switch(config- if)#`	Enters the configuration mode for the selected interface(s).
Step 3	`switch(config- if)# channel-group auto`	Automatically creates the channel group for the selected interface(s).
	`switch(config- if)# no channel-group auto`	Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled.

About Manually Configured Channel Groups

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

Tip If you enable persistence, be sure to enable it at both ends of the PortChannel.

Converting to Manually Configured Channel Groups

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

PortChannel Configuration Verification

You can view specific information about existing PortChannels at any time from EXEC mode. The following **show** commands provide further details on existing PortChannels. You can force all screen output to go to a printer or save it to a file. See Examples 16-1 to 16-6.

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

Example 16-1 Displays the PortChannel Summary

```
switch# show port-channel summary
--------------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
port-channel 77                2                 0                 --
port-channel 78                2                 0                 --
port-channel 79                2                 2                 fcip200
```

Example 16-2 Displays the PortChannel Configured in the Default ON Mode

```
switch# show port-channel database
port-channel 77
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
             fcip2    [down]
port-channel 78
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

Example 16-3 Displays the PortChannel Configured in the ACTIVE Mode

```
switch# show port-channel database
port-channel 77
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
```

```
            fcip2    [down]
port-channel 78
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

The **show port-channel consistency** command has two options—without and with details.

Example 16-4 Displays the Consistency Status without Details

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

Example 16-5 Displays the Consistency Status with Details

```
switch# show port-channel consistency detail
Authoritative port-channel database:
============================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
============================================
database 1: from module 5
============================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
============================================
database 2: from module 4
============================================
totally 3 port-channels
port-channel 77:
```

```
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
...
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

Example 16-6 Displays the PortChannel Usage

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

Example 16-7 Displays the PortChannel Compatibility

```
switch# show port-channel compatibility-parameters
physical port layer                 fibre channel or ethernet
    port mode                           E/AUTO only
    trunk mode
    speed
    port VSAN
    port allowed VSAN list
```

Use the existing **show** commands to obtain further details on autocreated channel group attributes. Autocreated PortChannels are indicated explicitly to help differentiate them from the manually created PortChannels. See Examples 16-8 to 16-10.

Example 16-8 Displays Autocreated PortChannels

```
switch# show interface fc1/1
fc1/1 is trunking
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:0a:00:0b:5f:3b:fe:80
    ...
    Receive data field Size is 2112
    Beacon is turned off
    Port-channel auto creation is enabled
    Belongs to port-channel 123
...
```

Example 16-9 Displays the Specified PortChannel Interface

```
switch# show port-channel database interface port-channel 128
port-channel 128
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    Channel is auto created
    First operational port is fc1/1
    1 ports in total, 1 ports up
    Ports:   fc1/1   [up] *
```

Example 16-10 Displays the PortChannel Summary

```
switch# show port-channel summary
--------------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
port-channel 1                 1                 0                 --
port-channel 2                 1                 1                 fc8/13
port-channel 3                 0                 0                 --
port-channel 4                 0                 0                 --
port-channel 5                 1                 1                 fc8/3
port-channel 6                 0                 0                 --
```

Default Settings

Table 16-3 lists the default settings for PortChannels.

Table 16-3 Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default PortChannel mode	ON.
Autocreation	Disabled.

CHAPTER 17

Configuring Domain Parameters

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

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

Tip 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.

This chapter includes the following sections:

Fibre Channel Domains

This section describes each fcdomain phase:

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

See Figure 17-1.

Figure 17-1 Sample fcdomain Configuration

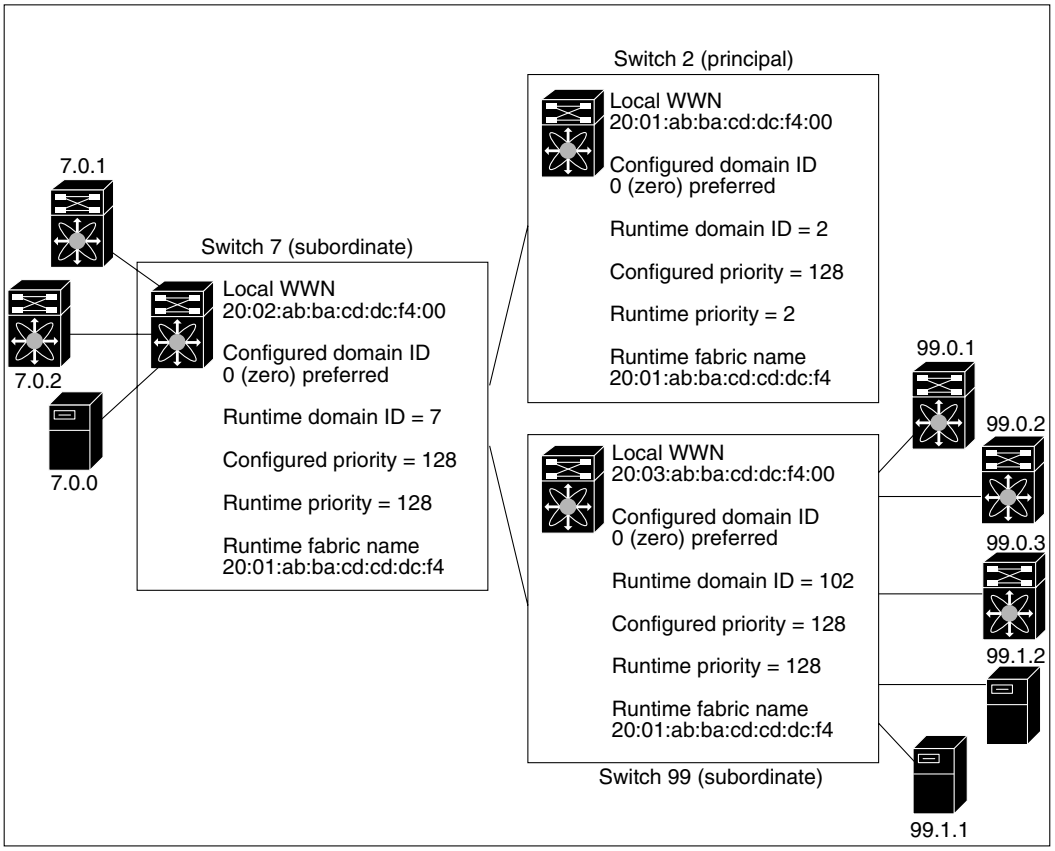

Note Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

This section describes the fcdomain feature and includes the following topics:

About Domain Restart

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

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

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

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

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 (see the "About Domain IDs" section on page 17-7).

Restarting a Domain

To restart the fabric disruptively or nondisruptively, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain restart vsan 1`	Forces the VSAN to reconfigure without traffic disruption.
	`switch(config)# fcdomain restart disruptive vsan 1`	Forces the VSAN to reconfigure with data traffic disruption.

About Domain Manager Fast Restart

As of Cisco MDS SAN-OS Release 3.0(2), 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.

Tip We recommend using fast restart on most fabrics, especially those with a large number of logical ports (3200 or more), where a logical port is an instance of a physical port in a VSAN.

Enabling Domain Manager Fast Restart

To enable the domain manager fast restart feature in Cisco SAN-OS Release 3.0(2) or later, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain optimize fast-restart vsan 3`	Enables domain manager fast restart on VSAN 3.
	`switch(config)# fcdomain optimize fast-restart vsan 7 - 10`	Enables domain manager fast restart on the range of VSANs from VSAN 7 to VSAN 10.
	`switch(config)# no fcdomain optimize fast-restart vsan 8`	Disables (default) domain manager fast restart on VSAN 8.

About Switch Priority

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

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

The priority configuration is applied to runtime when the fcdomain is restarted (see the "About Domain Restart" section on page 17-3). This configuration is applicable to both disruptive and nondisruptive restarts.

Configuring Switch Priority

To configure the priority for the principal switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain priority 25 VSAN 99`	Configures a priority of 25 for the local switch in VSAN 99.
	`switch(config)# no fcdomain priority 25 VSAN 99`	Reverts the priority to the factory default (128) in VSAN 99.

About fcdomain Initiation

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

Disabling or Reenabling fcdomains

To disable or reenable fcdomains in a single VSAN or a range of VSANs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no fcdomain vsan 7-200`	Disables the fcdomain configuration in VSAN 7 through 200.
	`switch(config)# fcdomain vsan 2008`	Enables the fcdomain configuration in VSAN 2008.

Configuring Fabric Names

To set the fabric name value for a disabled fcdomain, follow these steps:

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

About Incoming RCFs

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

The **rcf-reject** option takes immediate effect takes effect immediately. No fcdomain restart is required.

Rejecting Incoming RCFs

To reject incoming RCF request frames, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# fcdomain rcf-reject vsan 1`	Enables the RCF filter on the specified interface in VSAN 1.
	`switch(config-if)# no fcdomain rcf-reject vsan 1`	Disables (default) the RCF filter on the specified interface in VSAN 1.

About Autoreconfiguring Merged Fabrics

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

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

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

Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain auto-reconfigure vsan 10`	Enables the automatic reconfiguration option in VSAN 10.
	`switch(config)# no fcdomain auto-reconfigure 69`	Disables the automatic reconfiguration option and reverts it to the factory default in VSAN 69.

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.

This section describes how to configure domain IDs and includes the following topics:

About Domain IDs

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

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

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

When a subordinate switch requests a domain, the following process takes place (see Figure 17-2):

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

Figure 17-2 Configuration Process Using the preferred Option

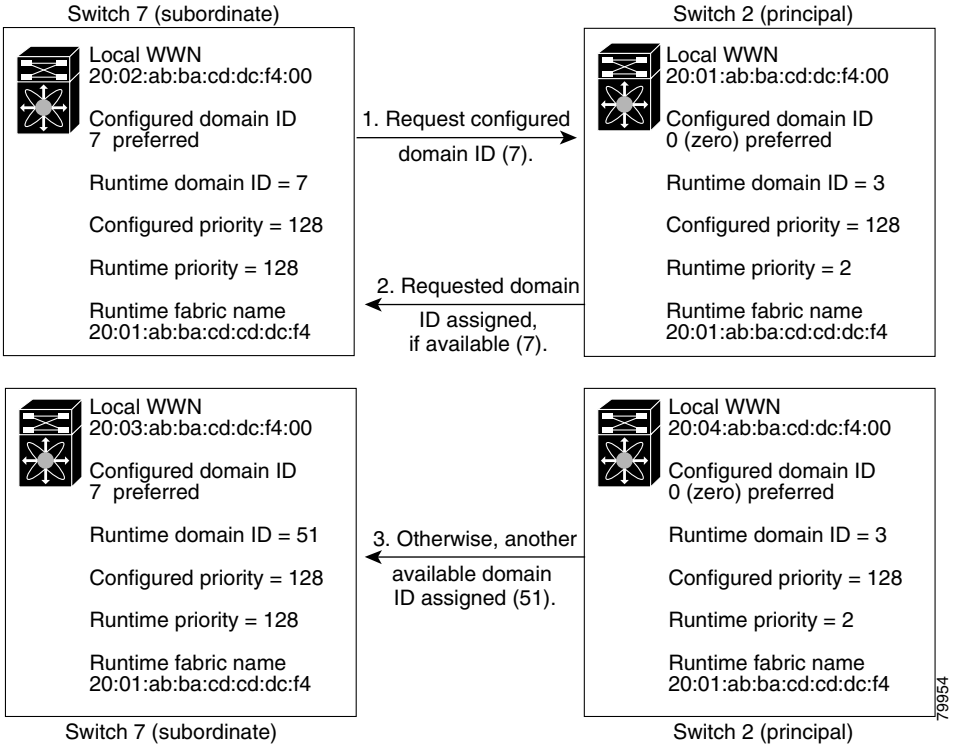

The behavior for a subordinate switch changes based on three factors:

- 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.

Tip When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN remains in the static state. You can change the static ID value but you cannot change it to the preferred option.

Note In an IVR without NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology should also be configured with static domain 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.

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

Note If you have configured an allow domain ID list, the domain IDs that you add must be in that range for the VSAN. See the "About Allowed Domain ID Lists" section on page 17-10.

Specifying Static or Preferred Domain IDs

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

While the static option can be applied at runtime after a disruptive or non-disruptive restart, the preferred option is applied at runtime only after a disruptive restart (see the "About Domain Restart" section on page 17-3).

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

To specify a static or preferred domain ID, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain domain 3 preferred vsan 8`	Configures the switch in VSAN 8 to request a preferred domain ID 3 and accepts any value assigned by the principal switch. The domain is range is 1 to 239.
	`switch(config)# no fcdomain domain 3 preferred vsan 8`	Resets the configured domain ID to 0 (default) in VSAN 8. The configured domain ID becomes 0 preferred.
Step 3	`switch(config)# fcdomain domain 2 static vsan 237`	Configures the switch in VSAN 237 to accept only a specific value and moves the local interfaces in VSAN 237 to an isolated state if the requested domain ID is not granted.
	`switch(config)# no fcdomain domain 18 static vsan 237`	Resets the configured domain ID to factory defaults in VSAN 237. The configured domain ID becomes 0 preferred.

About Allowed Domain ID Lists

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

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

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

An allowed domain ID list must satisfy the following conditions:

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

Configuring Allowed Domain ID Lists

To configure the allowed domain ID list, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain allowed 50-110 vsan 4`	Configures the list to allow switches with the domain ID 50 through 110 in VSAN 4.
	`switch(config)# no fcdomain allowed 50-110 vsan 5`	Reverts to the factory default of allowing domain IDs from 1 through 239 in VSAN 5.

To configure the allowed domain ID list using Fabric Manager, follow these steps:

Step 1 Expand **Fabric*xx* > VSAN*xx* > Domain Manager** and then select **Allowed** in the Logical Domains pane for the fabric and VSAN for which you want to set the allowed domain ID list.

About CFS Distribution of Allowed Domain ID Lists

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

Note All switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later to distribute the allowed domain ID list using CFS.

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

Note We recommend configuring the allow domain ID list and committing it on the principle switch.

For more information about CFS, see Chapter 6, “Using the CFS Infrastructure.”

Enabling Distribution

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

To enable (or disable) allowed domain ID list configuration distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain distribute`	Enables domain configuration distribution.
	`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. Once 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. Modifications from this point on are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

To apply the pending domain configuration changes to other MDS 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 MDS switches throughout the VSAN and the fabric lock is released.

To commit pending domain configuration changes and release the lock, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain commit vsan 10`	Commits the pending domain configuration changes.

Discarding Changes

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

To discard pending domain configuration changes and release the lock, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain abort vsan 10`	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.

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

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

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

Displaying CFS Distribution Status

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

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

Displaying Pending Changes

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

```
switch# show fcdomain pending vsan 10

Pending Configured Allowed Domains
---------------------------------

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

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

```
switch# show fcdomain pending-diff vsan 10

Current Configured Allowed Domains
---------------------------------

VSAN 10
Assigned or unallowed domain IDs: 24,100.
[User] configured allowed domain IDs: 1-239.

Pending Configured Allowed Domains
---------------------------------

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

Displaying Session Status

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

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

About Contiguous Domain ID Assignments

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

- 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 SAN-OS software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain contiguous-allocation vsan 81-83`	Enables the contiguous allocation option in VSAN 81 through 83. **Note** The **contiguous-allocation** option takes immediate effect at runtime. You do not need to restart the fcdomain.
	`switch(config)# no fcdomain contiguous-allocation vsan 1030`	Disables the contiguous allocation option and reverts it to the factory default in VSAN 1030.

FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following consequences apply:

- An N or NL port logs into a Cisco MDS 9000 Family switch. The WWN of the requesting N or NL 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.
- The switch connection behavior differs between N ports and NL ports:
  - 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).
  - NL ports receive the same FC IDs only if connected back to the same port on the switch to which they were originally connected.

This section describes configuring FC IDs and includes the following topics:

About Persistent FC IDs

When persistent FC IDs are enabled, the following consequences apply:

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

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

Note FC IDs are enabled by default. This change of default behavior from releases prior to Cisco MDS SAN-OS Release 2.0(1b) prevents FC IDs from being changed after a reboot. You can disable this option for each VSAN.

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

Note Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

Note Due to differences in Arbitrated Loop Physical Address (ALPA) support on devices, FC ID persistency for loop-attached devices is not guaranteed.

Enabling the Persistent FC ID Feature

To enable the persistent FC ID feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcdomain fcid persistent vsan 1000` `FCID(s) persistent feature is enabled.`	Activates (default) persistency of FC IDs in VSAN 1000.
	`switch(config)# no fcdomain fcid persistent vsan 20`	Disables the FC ID persistency feature in VSAN 20.

About Persistent FC ID Configuration

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. Follow these requirements to manually configure a persistent FC ID:

- 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.

Note FICON uses a different scheme for allocating FC IDs based in the front panel port number. This scheme takes precedence over FC ID persistence in FICON VSANs.

Configuring Persistent FC IDs

To configure persistent FC IDs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdomain fcid database` `switch(config-fcid-db)#`	Enters FC ID database configuration submode.
Step 3	`switch(config-fcid-db)# vsan 1000 wwn 33:e8:00:05:30:00:16:df fcid 0x070128`	Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in VSAN 1000. **Note** To avoid assigning a duplicate FC ID, use the **show fcdomain address-allocation vsan** command to display the FC IDs in use.
	`switch(config-fcid-db)# vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070123 dynamic`	Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in VSAN 1000 in dynamic mode.
	`switch(config-fcid-db)# vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070100 area`	Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x701FF in VSAN 1000. **Note** To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID.

About Unique Area FC IDs for HBAs

Note Only read this section if the HBA port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than 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.

Switches in the Cisco MDS 9000 Family 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. The procedure in this example uses a switch domain of 111(6f hex). The HBA port connects to interface fc1/9 and the storage port connects to interface fc 1/10 in the same switch.

Configuring Unique Area FC IDs for an HBA

To configure a different area ID for the HBA port, follow these steps:

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

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE    VSAN    FCID        PORT NAME                NODE NAME
--------------------------------------------------------------------------------
 fc1/9       3       0x6f7703    50:05:08:b2:00:71:c8:c2  50:05:08:b2:00:71:c8:c0
 fc1/10      3       0x6f7704    50:06:0e:80:03:29:61:0f  50:06:0e:80:03:29:61:0f
```

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

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

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# shutdown
switch(config-if)# end
switch#
```

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 with this procedure to enable the persistent FC ID.

If this feature is already enabled, skip to Step 5.

Step 4 Enable the persistent FC ID feature in the Cisco MDS switch.

```
switch# conf t
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
switch#
```

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

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

Step 6 Enable the HBA interface in the Cisco MDS switch.

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# no shutdown
switch(config-if)# end
switch#
```

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

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE    VSAN    FCID        PORT NAME                 NODE NAME
--------------------------------------------------------------------------------
fc1/9        3       0x6fee00    50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
fc1/10       3       0x6f7704    50:06:0e:80:03:29:61:0f   50:06:0e:80:03:29:61:0f
```

Note Both FC IDs now have different area assignments.

About Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. Table 17-1 identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 17-1 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

To purge persistent FC IDs, follow this step:

	Command	Purpose
Step 1	`switch# purge fcdomain fcid vsan 4`	Purges all dynamic and unused FC IDs in VSAN 4.
	`switch# purge fcdomain fcid vsan 3-5`	Purges dynamic and unused FC IDs in VSAN 3, 4, and 5.

Displaying fcdomain Information

Use the **show fcdomain** command to display global information about fcdomain configurations. See Example 17-1.

Note In Example 17-1, the fcdomain feature is disabled. Consequently, the runtime fabric name is the same as the configured fabric name.

Example 17-1 Displays the Global fcdomain Information

```
switch# show fcdomain vsan 2
The local switch is the Principal Switch.

Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:0b:46:79:ef:41
        Running fabric name: 20:01:00:0b:46:79:ef:41
        Running priority: 128
        Current domain ID: 0xed(237)

Local switch configuration information:
        State: Enabled
        FCID persistence: Disabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)

Principal switch run time information:
        Running priority: 128

No interfaces available.
```

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. Example 17-2 shows the following:

- 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.

Example 17-2 Displays the fcdomain Lists

```
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. See Example 17-3.

Example 17-3 Displays the Allowed Domain ID Lists

```
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.
```

Tip Ensure that the requested domain ID passes the Cisco SAN-OS software checks, if **interop 1** mode is required in this switch.

Use the **show fcdomain fcid persistent** command 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. See Examples 17-4 and 17-5.

Example 17-4 Displays Persistent FC IDs in a Specified VSAN

```
switch# show fcdomain fcid persistent vsan 1000
Total entries 2.

Persistent FCIDs table contents:
VSAN              WWN                FCID         Mask         Used    Assignment
----    ----------------------    --------    -----------    ----    ----------
1000    11:11:22:22:11:11:12:23    0x700101    SINGLE FCID      NO    STATIC
1000    44:44:33:33:22:22:11:11    0x701000    ENTIRE AREA      NO    DYNAMIC
```

Example 17-5 Displays All Persistent FC IDs in the fcdomain

```
switch# show fcdomain fcid persistent
Total entries 2.

Persistent FCIDs table contents:
VSAN              WWN                FCID         Mask         Used    Assignment
----    ----------------------    --------    -----------    ----    ----------
1000    11:11:22:22:11:11:22:22    0x700501    SINGLE FCID      NO    STATIC
1003    44:44:33:33:22:22:11:11    0x781000    ENTIRE AREA     YES    DYNAMIC
```

Use the **show fcdomain statistics** command to display frame and other fcdomain statistics for a specified VSAN or PortChannel. See Example 17-6 and Example 17-7.

Example 17-6 Displays fcdomain Statistics for a Specified VSAN

```
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
```

Example 17-7 Displays fcdomain Statistics for a Specified PortChannel

```
switch# show fcdomain statistics interface port-channel 10 vsan 1
        Interface Statistics:
                Transmitted      Received
                -----------      --------
                EFPs      13       9
                DIAs      7        7
                RDIs      0        0
                ACCs      21       25
                RJTs      1        1
                BFs       2        2
                RCFs      4        4
                Error     0        0
                Total     48       48
        Total Retries: 0
        Total Frames: 96
                -----------      --------
```

Use the **show fcdomain address-allocation** command to display FC ID allocation statistics including a list of assigned and free FC IDs. See Example 17-8.

Example 17-8 Displays FC ID Information

```
switch# show fcdomain address-allocation vsan 1
Free FCIDs: 0x020000 to 0x02fdff
            0x02ff00 to 0x02fffe

Assigned FCIDs: 0x02fe00 to 0x02feff
                0x02ffff

Reserved FCIDs: 0x020100 to 0x02f0ff
                0x02fe00 to 0x02feff
                0x02ffff


Number free FCIDs: 65279
Number assigned FCIDs: 257
Number reserved FCIDs: 61697
```

Use the **show fcdomain address-allocation cache** command 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. See Example 17-9.

Example 17-9 Displays Address Allocation Information

```
switch# show fcdomain address-allocation cache
Cache content:
line#    VSAN              WWN                FCID         mask
-----    ----    ----------------------    --------    -----------
   1.      12    21:00:00:e0:8b:08:a2:21    0xef0400    ENTIRE AREA
   2.       6    50:06:04:82:c3:a1:2f:5c    0xef0002    SINGLE FCID
   3.       8    20:4e:00:05:30:00:24:5e    0xef0300    ENTIRE AREA
   4.       8    50:06:04:82:c3:a1:2f:52    0xef0001    SINGLE FCID
```

Default Settings

Table 17-2 lists the default settings for all fcdomain parameters.

Table 17-2 Default fcdomain Parameters

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

CHAPTER 18

Scheduling Maintenance Jobs

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family.You can use this feature to schedule jobs on a one-time basis or periodically.

This chapter includes the following sections:

About the Command Scheduler

The Cisco SAN-OS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future. The job(s) can be executed once at a specified time in the future or at periodic intervals.

Note To use the command scheduler, you do not need to obtain any license.

You can use this feature to schedule zone set changes, QOS policy changes, backup data, save the configuration and other similar jobs.

Scheduler Terminology

The following terms are used in this chapter.

- Job—A job is a set of SAN-OS CLI commands (EXEC and config mode) that are executed as defined in the schedule.
- Schedule—A schedule determines the time when the assigned jobs must be executed. Multiple jobs can be assigned to a schedule. A schedule executes in one of two modes: one-time or periodic.
- Periodic mode—A job is executed at the user-specified periodic intervals, until it is deleted by the administrator. The following types of periodic intervals are supported:
  - Daily—The job is executed once a day.
  - Weekly—The job is executed once a week.

  - Monthly—The job is executed once a month.
  - Delta—The job is executed beginning at the specified start time and thereafter at user-specified intervals (days:hours:minutes).
- One-time mode—The job is executed once at a user-specified time.

Scheduling Guidelines

Before scheduling jobs on a Cisco MDS switch, be aware of the following guidelines:

- Prior to Cisco MDS SAN-OS Release 3.0(3), only users local to the switch could perform scheduler configuration. As of Cisco MDS SAN-OS Release 3.0(3), remote users can perform job scheduling using AAA authentication.
- Be aware that the scheduled job can fail if it encounters one of the following situations when executing the job:
  - If the license has expired for a feature at the time when a job containing commands pertaining to that feature is scheduled.
  - If a feature is disabled at the time when a job containing commands pertaining to that feature is scheduled.
  - If you have removed a module from a slot and the job has commands pertaining to the interfaces for that module or slot.
- Verify that you have configured the time. The scheduler does not have any default time configured. If you create a schedule and assign job(s) and do not configure the time, that schedule is not launched.
- While defining a job, verify that no interactive or disruptive commands (for example, **copy bootflash:** *file* **ftp:** *URI*, **write erase**, and other similar commands) are specified as part of a job because the job is executed noninteractively at the scheduled time.

Configuring the Command Scheduler

To configure the command scheduler, follow these steps:

Step 1 Enable the scheduler.

Step 2 Authorize remote user access (optional).

Step 3 Define the job.

Step 4 Specify the schedule and assign jobs to the schedule.

Step 5 Specify the time for the schedule(s).

Step 6 Verify the scheduled configuration.

This section includes the following topics:

Enabling the Command Scheduler

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

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

To enable the command scheduling feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# scheduler enable`	Enables the command scheduler.
	`switch(config)# no scheduler enable`	Discards the scheduler configuration and disables the command scheduler (default).

To display the command schedule status, use the **show scheduler config** command.

```
switch# show scheduler config
config terminal
  scheduler enable
  scheduler logfile size 16
end
```

Configuring Remote User Authentication

Prior to Cisco MDS SAN-OS Release 3.0(3), only users local to the switch could perform scheduler configuration. As of Cisco MDS SAN-OS Release 3.0(3), remote users can perform job scheduling using AAA authentication.

Note AAA authentication requires the clear text password of the remote user before creating and configuring command scheduler jobs.

To configure remote user authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# scheduler aaa-authentication password X12y34Z56a`	Configures a clear text password for remote users.
Step 3	`switch(config)# scheduler aaa-authentication password 0 X12y34Z56a`	Configures a clear text password for remote users.
Step 4	`switch(config)# no scheduler aaa-authentication password`	Removes the clear text password for remote users.
Step 5	`switch(config)#scheduler aaa-authentication user newuser password Z98y76X54b`	Configures a clear text password for remote user newuser.
Step 6	`switch(config)#scheduler aaa-authentication user newuser password 0 Z98y76X54b`	Configures a clear text password for remote user newuser.
Step 7	`switch(config)# no scheduler aaa-authentication password user newuser`	Removes the clear text password for remote user newuser.

To display the scheduler password configuration for remote users, use the **show running-config** command.

```
switch# show running-config | include "scheduler aaa-authentication"
scheduler aaa-authentication username newuser password 7 "C98d76S54e"
```

Note The scheduler remote user passwords are always displayed in encrypted form in the **show running-config** command output. The encrypted option (**7**) in the command exists to support applying the ASCII configuration to the switch.

Defining a Job

To define a job, you must specify the job name. This action places you in the job definition (`config-job`) submode. In this submode, you can define the sequence of CLI commands that the job has to perform. Be sure to exit the `config-job` submode to complete the job definition.

Caution You cannot modify or remove a command after entering the sequence of commands. To make changes, you must explicitly delete the defined job name and restart this process.

Note You must exit the config-job submode for the job definition to be complete.

To define a job for the command scheduler, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler job name addMemVsan99` `switch(config-job)#`	Defines a job name and enters the job definition submode

Step 3	`switch(config-job)# config terminal`
	`switch(config-job-config)# vsan database`
	`switch(config-job-config-vsan-db)# vsan 99 interface fc1/1 - 4`
	`switch(config-job-config-vsan-db)# end`
	`switch#`
		Specifies a sequence of actions for the specified job. The defined commands are checked for validity and stored for future use.
		Note Be sure you exit the `config-job` submode.

	`switch(config)# scheduler job name offpeakQOS`
	`switch(config-job)# conf t`
	`switch(config-job-config)# qos class-map offpeakbackupcmap match-all`
	`switch(config-job-config-cmap)# match source-wwn 23:15:00:05:30:00:2a:1f`
	`switch(config-job-config-cmap)# match destination-wwn 20:01:00:05:30:00:28:df`
	`switch(config-job-config-cmap)# exit`
	`switch(config-job-config)# qos policy-map offpeakbackuppolicy`
	`switch(config-job-config-pmap)# class offpeakbackupcmap`
	`switch(config-job-config-pmap-c)# priority high`
	`switch(config-job-config-pmap-c)# exit`
	`switch(config-job-config-pmap)# exit`
	`switch(config-job-config)# qos service policy offpeakbackuppolicy vsan 1`
	`switch(config-job-config)# end`
	`switch#`
		Provides another example of scheduling a different set of jobs.

Verifying the Job Definition

To verify the job definition, use the **show scheduler job** command.

```
switch# show scheduler job addMemVsan99
Job Name: addMemVsan99
----------------------
  config terminal
   vsan database
    vsan 99 interface fc1/1
    vsan 99 interface fc1/2
    vsan 99 interface fc1/3
    vsan 99 interface fc1/4
```

Deleting a Job

To delete a job for the command scheduler, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# no scheduler job name addMemVsan99`	Deletes a defined job and all commands defined within that job.

Specifying a Schedule

After defining jobs, you can create schedules and assign jobs to the schedule. Subsequently, you can configure the time of execution. The execution can be one-time or periodic depending on your requirements. If the time for the schedule is not configured, then it will never be executed.

Specifying a Periodic Schedule

When you specify a periodic job execution, that job is executed periodically at the specified (daily, weekly, monthly, or delta) intervals.

To specify a periodic job for the command scheduler, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler schedule name weekendbackupqos` `switch(config-schedule)#`	Defines a job schedule (weekendbackup) and enters the submode for that schedule.
	`switch(config)# no scheduler schedule name weekendbackup`	Deletes the defined schedule.
Step 3	`switch(config-schedule)# job name offpeakZoning` `switch(config-schedule)# job name offpeakQOS`	Assign two jobs offpeakZoning and offpeakQOS for this schedule.
Step 4	`switch(config-schedule)# no job name addMem99`	Deletes the job assigned for this schedule.
Note	The following examples are only provided for reference.	

	Command	Purpose
Step 5	`switch(config-schedule)# time daily 23:00`	Executes the specified jobs at 11 p.m. every day.
	`switch(config-schedule)# time weekly Sun:23:00`	Specifies a weekly execution every Sunday at 11 p.m.
	`switch(config-schedule)# time monthly 28:23:00`	Specifies a monthly execution at 11 p.m on the 28th of each month. If you specify the date as either 29, 30, or 31, the command is automatically executed on the last day of each month.
	`switch(config-schedule)# time start now repeat 48:00`	Specifies a job to be executed every 48 hours beginning 2 minutes from *now* —if today is September 24, 2004, and the time is now 2:00 p.m., the command begins executing at 2 minutes past 2:00 p.m. on September 24, 2004, and continues to execute every 48 hours after that.
	`switch(config-schedule)# time start 14:00 repeat 14:00:00`	If today is September 24, 2004, (Friday), this command specifies the job to be executed every alternate Friday at 2 p.m. (every 14 days).

The most significant fields in the **time** parameter are optional. If you omit the most significant fields, the values are assumed to be the same as the current time. For example, if the current time is September 24, 2004, 22:00 hours, then the commands are executed as follows:

- The **time start 23:00 repeat 4:00:00** command implies a start time of September 24, 2004, 23:00 hours.
- The **time daily 55** command implies every day at 22:55 hours.
- The **time weekly 23:00** command implies every Friday at 23:00 hours.
- The **time monthly 23:00** command implies the 24th of every month at 23:00 hours.

Note If the time interval configured for any schedule is smaller than the time taken to execute its assigned job(s), then the subsequent schedule execution occurs only after the configured interval amount of time has elapsed following the completion time of the last iteration of the schedule. For example, a schedule is executed at 1-minute intervals and a job assigned to it takes 2 minutes to complete. If the first schedule is at 22:00 hours, the job finishes at 22:02 after which, the 1-minute interval is observed and the next execution occurs at 22:03 and finishes at 22:05.

Specifying a One-Time Schedule

When you specify a one-time job execution, that job is only executed once

To specify a one-time job for the command scheduler, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler schedule name configureVsan99` `switch(config-schedule)#`	Defines a job schedule (configureVsan99) and enters the submode for that schedule.

	Command	Purpose
Step 3	`switch(config-schedule)# job name addMemVsan99`	Assigns a predefined job name (addMemVsan99) for this schedule.
Step 4	`switch(config-schedule)# time start 2004:12:14:23:00`	Specifies a one-time execution on December 14, 2004, at 11 p.m.
	`switch(config-schedule)# no time`	Deletes the time assigned for this schedule.

Verifying Scheduler Configuration

To display the scheduler configuration, use the **show scheduler config** command.

```
switch# show scheduler config
config terminal
  scheduler enable
  scheduler logfile size 512
end

config terminal
   scheduler job name addMemVsan99
     config terminal
       vsan database
       vsan 99 interface fc1/1
       vsan 99 interface fc1/2
       vsan 99 interface fc1/3
       vsan 99 interface fc1/4
end

config terminal
  scheduler schedule name configureVsan99
    time start 2004:8:10:9:52
    job name addMemVsan99
end
```

Deleting a Schedule

To delete a schedule, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# no scheduler schedule name weekendbackup`	Deletes the defined schedule.

Removing an Assigned Job

To remove an assigned job, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler schedule name weekendbackupqos` `switch(config-schedule)#`	Specifies a job schedule (weekendbackupqos) and enters the submode for that schedule.
Step 3	`switch(config-schedule)# no job name addMem99`	Removes a job (addMem99) assigned to this schedule.

Deleting a Schedule Time

To delete the schedule time, follow these steps:

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler schedule name weekendbackupqos` `switch(config-schedule)#`	Defines a job schedule (weekendbackup) and enters the submode for that schedule.
Step 3	`switch(config-schedule)# no time`	Deletes the schedule time configuration.The schedule will not be run until the time is configured again.

Verifying the Command Scheduler Execution Status

To verify the command scheduler execution status, use the **show scheduler schedule** command.

```
switch# show scheduler schedule configureVsan99
Schedule Name       : configureVsan99
-----------------------------------
User Name           : admin
Schedule Type       : Run once on Tue Aug 10 09:48:00 2004
Last Execution Time: Tue Aug 10 09:48:00 2004
-----------------------------------------------
          Job Name                  Status
-----------------------------------------------
addMemVsan99                        Success (0)
```

Execution Logs

This section describes execution logs for the command scheduler and contains the following sections:

About Execution Logs

The command scheduler maintains a log file. While you cannot modify the contents of this file, you can change the file size. This log file is a circular log that contains the output of the job executed. If the output of the job is greater than the log file, then the output stored in this file remains truncated.

You can configure the log file size to be a maximum of 1024 KB. The default size of the execution log file is 16 KB.

Configuring Execution Logs

To configure the execution log file size, follow these steps:\

	Command	Purpose
Step 1	`switch# conf t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# scheduler logfile size 1024`	Configures the log file to be a maximum of 1024 KB
	`switch(config)# no scheduler logfile size`	Defaults to the log size of 16 KB.

To display the execution log file configuration, use the **show scheduler config** command.

```
switch# show scheduler config
config terminal
  scheduler enable
  scheduler logfile size 1024
end
```

Displaying Execution Log File Contents

To display the execution log for all jobs executed in the system, use the **show scheduler logfile** command.

```
switch# show scheduler logfile
Job Name       : addMemVsan99              Job Status: Success (0)
Schedule Name  : configureVsan99           User Name : admin
Completion time: Tue Aug 10 09:48:00 2004
-------------------------- Job Output --------------------------
`config terminal`
`vsan database`
`vsan 99 interface fc1/1`
`vsan 99 interface fc1/2`
`vsan 99 interface fc1/3`
`vsan 99 interface fc1/4`
```

Clearing the Execution Log File Contents

To clear the contents of the scheduler execution log file, issue the **clear scheduler logfile** command in EXEC mode.

```
switch# clear scheduler logfile
```

Default Settings

Table 18-1 lists the default settings for command scheduling parameters.

Table 18-1 Default Command Scheduler Parameters

Parameters	Default
Command scheduler	Disabled.
Log file size	16 KB.

PART 4

Fabric Configuration

CHAPTER 19

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. This chapter includes the following sections:

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.

This section describes VSANs and includes the following topics:

VSANs Topologies

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

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.

- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

The switch icons shown in both Figure 19-1 and Figure 19-2 indicate that these features apply to any switch in the Cisco MDS 9000 Family.

Figure 19-1 shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. No communication between VSANs is possible. Within each VSAN, all members can talk to one another.

Figure 19-1 Logical VSAN Segmentation

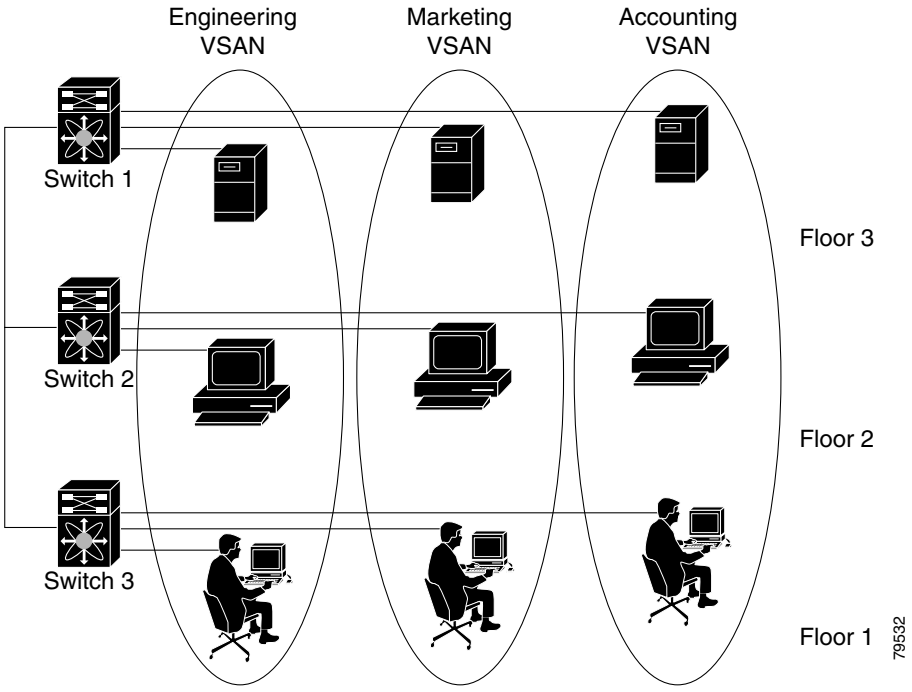

Figure 19-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Figure 19-2 Example of two VSANs

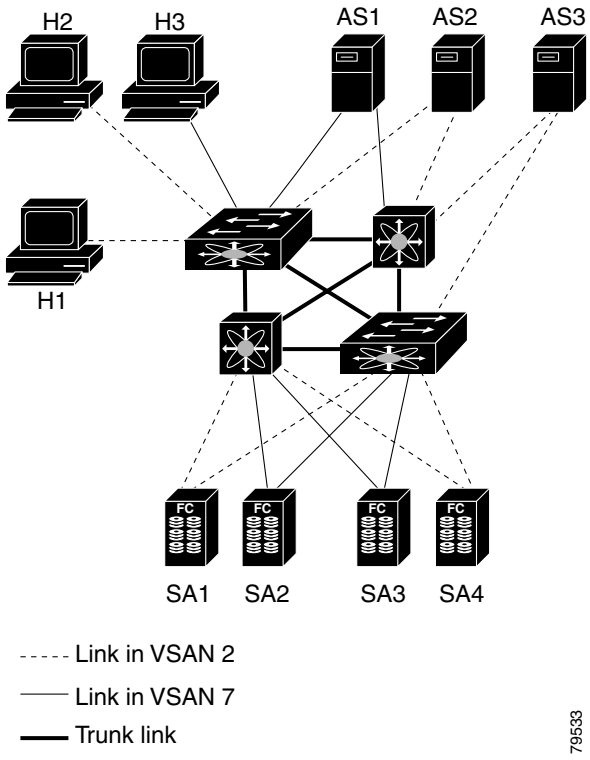

The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. Thus the inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN topologies.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. Figure 19-2 illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

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

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

VSAN Advantages

VSANs offer the following advantages:

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

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

VSANs Versus Zones

You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 19-1 lists the differences between VSANs and zones.

Table 19-1 VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
—	Zones are always contained within a VSAN. Zones never span two VSANs.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to Fx ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN—the VSAN associated with the Fx port.	An HBA or storage device can belong to multiple zones.
VSANs enforce membership at each E port, source port, and destination port.	Zones enforce membership only at the source and destination ports.
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

Figure 19-3 shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 19-3 VSANS with Zoning

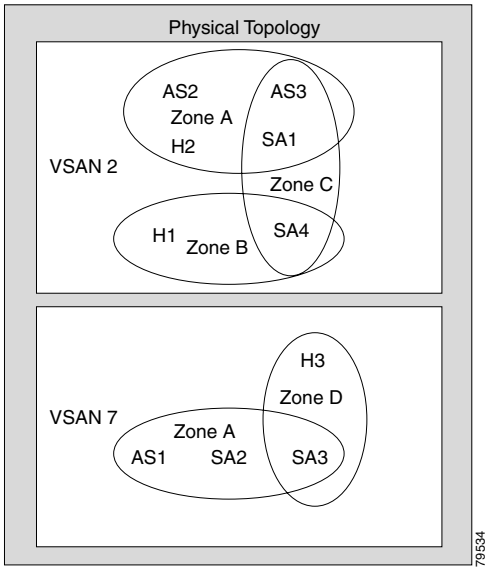

VSAN Configuration

VSANs have the following attributes:

- VSAN ID—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- State—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
  - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
  - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.

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

Note A VSAN name must be unique.

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

Note OX ID based load balancing of IVR traffic from IVR- enabled switches is not supported on Generation 1 switching modules. OX ID based load balancing of IVR traffic from a non-IVR MDS switch should work. Generation 2 switching modules support OX ID based load balancing of IVR traffic from IVR-enabled switches.

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

About VSAN Creation

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

Creating VSANs Statically

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

To create VSANs, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch(config-vsan-db)#`	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.
Step 3	`switch(config-vsan-db)# vsan 2`	Creates a VSAN with the specified ID (2) if that VSAN does not exist already.
Step 4	`switch(config-vsan-db)# vsan 2 name TechDoc` `updated vsan 2`	Updates the VSAN with the assigned name (TechDoc).
Step 5	`switch(config-vsan-db)# vsan 2 suspend`	Suspends the selected VSAN.
Step 6	`switch(config-vsan-db)# no vsan 2 suspend`	Negates the **suspend** command issued in the previous step.
Step 7	`switch(config-vsan-db)# end` `switch#`	Returns you to EXEC mode.

About Port VSAN Membership

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

- Statically—by assigning VSANs to ports.

  See the "Assigning Static Port VSAN Membership" section on page 19-7.

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

  See Chapter 21, "Creating Dynamic VSANs."

Trunking ports have an associated list of VSANs that are part of an allowed list (see Chapter 15, "Configuring Trunking").

Assigning Static Port VSAN Membership

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch(config-vsan-db)#`	Configures the database for a VSAN.
Step 3	`switch(config-vsan-db)# vsan 2`	Creates a VSAN with the specified ID (2) if that VSAN does not exist already.
Step 4	`switch(config-vsan-db)# vsan 2 interface fc1/8`	Assigns the membership of the fc1/8 interface to the specified VSAN (VSAN 2).
Step 5	`switch(config-vsan-db)# vsan 7`	Creates another VSAN with the specified ID (7) if that VSAN does not exist already.

	Command	Purpose
Step 6	`switch(config-vsan-db)# vsan 7 interface fc1/8`	Updates the membership information of the interface to reflect the changed VSAN.
	`switch(config-vsan-db)# no vsan 7 interface fc1/8`	Removes the interface from the VSAN.

Displaying VSAN Static Membership

To display the VSAN static membership information, use the **show vsan membership** command (see Example 19-1 through Example 19-3).

Example 19-1 Displays Membership Information for the Specified VSAN

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

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

Example 19-2 Displays Static Membership Information for All VSANs

```
switch # show vsan membership
vsan 1 interfaces:
        fc2/16  fc2/15  fc2/14  fc2/13  fc2/12  fc2/11  fc2/10  fc2/9
        fc2/8   fc2/7   fc2/6   fc2/5   fc2/4   fc2/3   fc2/2   fc2/1
        fc1/16  fc1/15  fc1/14  fc1/13  fc1/12  fc1/11  fc1/10  fc1/9
        fc1/7   fc1/6   fc1/5   fc1/4   fc1/3   fc1/2   fc1/1
vsan 2 interfaces:
        fc1/8
vsan 7 interfaces:
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

Example 19-3 Displays Static Membership Information for a Specified Interface

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

About the Default VSAN

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. We recommend that you do not use VSAN 1 as your production environment VSAN. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.

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

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

About the Isolated VSAN

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).

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

Caution Do not use an isolated VSAN to configure ports.

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

Displaying Isolated VSAN Membership

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

Operational State of a VSAN

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

About Static VSAN Deletion

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

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see Figure 19-4).

Figure 19-4 VSAN Port Membership Details

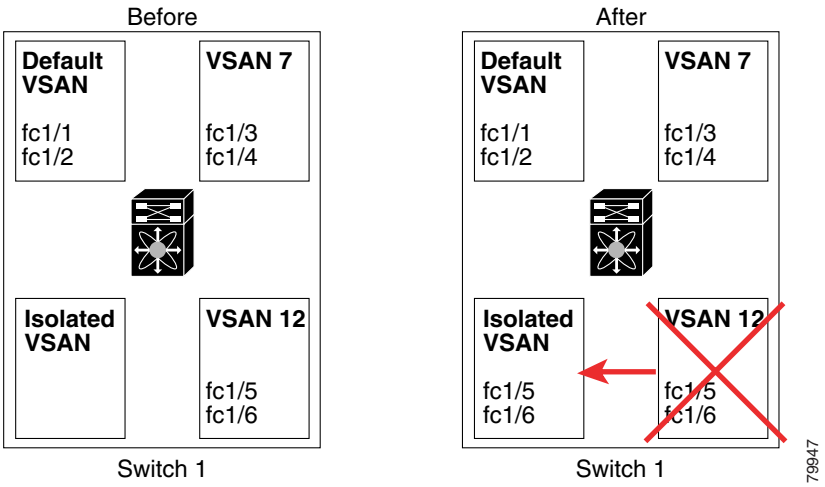

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

Note The allowed VSAN list is not affected when a VSAN is deleted (see Chapter 15, "Configuring Trunking").

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

Deleting Static VSANs

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch(config-db)#`	Configures the VSAN database.
Step 3	`switch-config-db# vsan 2` `switch(config-vsan-db)#`	Places you in VSAN configuration mode.

	Command	Purpose
Step 4	`switch(config-vsan-db)# no vsan 5` `switch(config-vsan-db)#`	Deletes VSAN 5 from the database and switch.
Step 5	`switch(config-vsan-db)# end` `switch#`	Places you in EXEC mode.

About Load Balancing

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

Configuring Load Balancing

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch(config-vsan-db)#`	Enters VSAN database configuration submode
Step 3	`switch(config-vsan-db)# vsan 2`	Specifies an existing VSAN.
Step 4	`switch(config-vsan-db)# vsan 2 loadbalancing src-dst-id`	Enables the load balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process.
	`switch(config-vsan-db)# no vsan 2 loadbalancing src-dst-id`	Negates the command issued in the previous step and reverts to the default values of the load balancing parameters.
	`switch(config-vsan-db)# vsan 2 loadbalancing src-dst-ox-id`	Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default).
Step 5	`switch(config-vsan-db)# vsan 2 suspend`	Suspends the selected VSAN.
Step 6	`switch(config-vsan-db)# no vsan 2 suspend`	Negates the **suspend** command issued in the previous step.
Step 7	`switch(config-vsan-db)# end` `switch#`	Returns you to EXEC mode.

About Interop Mode

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces. See the "Switch Interoperability" section on page 29-11.

About FICON VSANs

You can enable FICON in up to eight VSANs. See the "FICON VSAN Prerequisites" section on page 28-7.

Displaying Static VSAN Configuration

Use the **show vsan** command to display information about configured VSANs (see Examples 19-4 to 19-6).

Example 19-4 Displays the Configuration for a Specific VSAN

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

Example 19-5 Displays the VSAN Usage

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

Example 19-6 Displays All VSANs

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

Default Settings

Table 19-2 lists the default settings for all configured VSANs.

Table 19-2 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 20

SAN Device Virtualization

This chapter describes how to configure virtual devices to represent physical end devices for switches running Cisco MDS SAN-OS Release 3.1(2) and later.

Cisco SAN device virtualization (SDV) is a licensed feature included in the Cisco MDS 9000 Family Enterprise package (ENTERPRISE_PKG). See Chapter 3, "Obtaining and Installing Licenses," for details about acquiring licenses.

This chapter includes the following sections:

About SDV

As of Cisco SAN-OS Release 3.1(2) and later, you can use Cisco SDV to create virtual devices that represent physical end-devices. Virtualization of SAN devices accelerates swapout or failover to a replacement disk array, and it also minimizes downtime when replacing host bus adapters (HBAs) or when re-hosting an application on a different server.

SAN devices that are virtualized can be either initiators or targets. You can virtualize targets to create a *virtual target*, and also virtualize initiators to create a *virtual initiator*. Such configurations do not distinguish between virtual initiators and virtual targets (see Figure 20-1 and Figure 20-2).

Figure 20-1 Target Virtualization

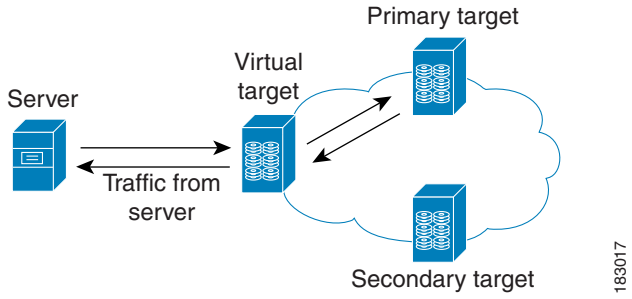

Figure 20-2 Initiator Virtualization

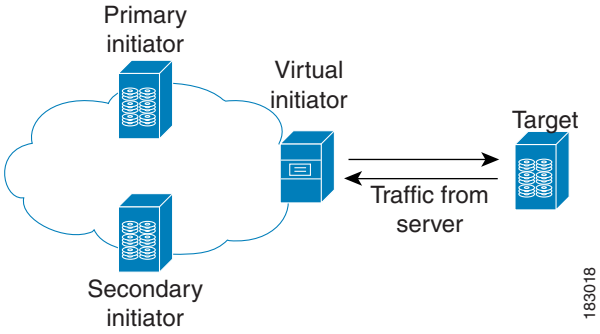

Note While most of the examples in this chapter describe target virtualization, the same behaviors apply to initiator virtualization as well.

Typically, today's deployments for handling device failures are designed for high availability (HA), with redundancy being a key part of this design. Let's consider the case where a target is designed to be redundant. Here, two arrays are deployed–a primary and secondary. Enterprises often use some type of consistency technology (such as EMF SRDF) between the primary and secondary arrays to ensure that the secondary is a mirrored copy of the production LUN. However, if the primary array fails, it must be replaced by the secondary, as all I/O must occur on the secondary. Problems can occur because the time required to bring the secondary array up and have it working often takes longer than most can afford (Figure 20-3 illustrates this dilemma).

Figure 20-3 Typical Deployment for Handling Device Failures Before SDV

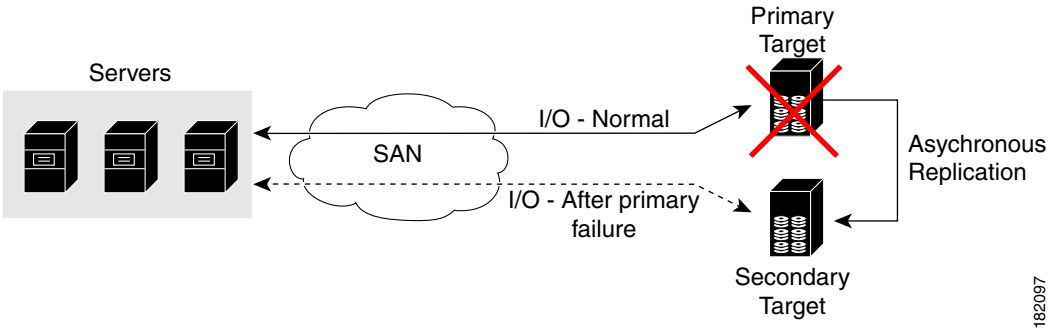

If a storage array is replaced *without* using Cisco SDV, then it may require the following:

- Taking down a server to modify zoning and account for the new array.
- Changing the Cisco SAN-OS configuration to accommodate Fibre Channel IDs (FC IDs) and pWWNs of the new array.
- Changing a server configuration to accommodate the new FC IDs and pWWNs.

More specifically, without SDV you might experience the following:

- It can take a considerable amount of time to configure a secondary device for a typical production environment.
- In the zoning configuration, all the initiators must be re-zoned with the secondary device, and certain initiators must also be reconfigured. For example, the WWN and FC ID of the secondary device are different, so driver files must be changed and the server must be rebooted.
- Clustering (multiple initiators) compounds the problem, and the failover procedure must be repeated for each server of the cluster. Think of a server cluster as a set of HBAs–any storage array FC ID changes must be performed for each HBA.

SDV enables you to:

- Reduce the amount of time it takes for data migration, and ultimately the overall amount of downtime.
- Easily scale to larger numbers of devices.

Figure 20-4 illustrates the benefits of SDV. In this configuration, disk array Y replaces disk array X. When disk array X was deployed, the user created virtual devices for all the Fibre Channel interfaces using SDV. After data replication from disk array X was completed, the user briefly pauses activity on the application server and re-linked disk array Y to the virtual devices used by the server, completing the swapout of disk array X. No zoning changes or host operating system configuration changes were required during the time-critical period when the swap was performed; this significantly minimized application downtime.

Note The array administrator will likely have to perform actions on array Y for it to become a primary device and accept server log ins before linking the virtual device to the array Y pWWN.

Figure 20-4 SDV Example

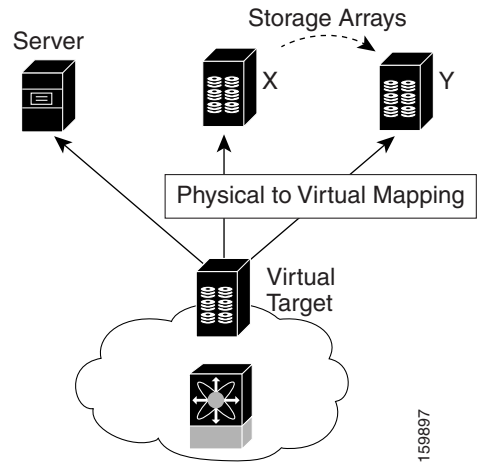

Key Concepts

The following terms are used throughout this chapter:

- Virtual device
  The virtualized or proxy representation of the real device, which is registered with the name server and has a pWWN and FC ID. A virtual device exists as long as its real (physical) counterpart is online. The virtual device pWWN and FC ID must be unique and cannot clash with any real device pWWNs and FC IDs.
- Virtual domain
  Reserved by SDV to assign FC IDs to virtual devices. If the switch that reserved the domain goes down, another switch takes over its role using the same domain.

Configuring SDV

SDV is a distributed service and uses CFS (Cisco Fabric Services) distribution to synchronize the databases. When you configure SDV it starts a CFS session and locks the fabric. When a fabric is locked, Cisco SAN-OS software does not allow any configuration changes from a switch–other than the switch holding the lock–and issues a message to inform users about the locked status. Configuration changes are held in a pending database for the application. You must perform a commit operation to make the configuration active and to release the lock for all switches. You can discard or stop changes from being distributed by issuing the **abort/clear** command.

See Chapter 6, “Using the CFS Infrastructure” for more details about CFS,

Note When you enable SDV, CFS distribution is also enabled; CFS distribution cannot be disabled for SDV.

The following sections describe how to configure SDV:

Configuring a Virtual Device

A virtual device is identified by an alphanumeric name of up to 32 characters and defines all the real devices (one primary and one or more secondary) that it represents. Upon the successful creation of a virtual device, the virtual device name is internally registered as the device alias name with the device alias database; the pWWN is automatically assigned by the system using Cisco OUI (Organizational Unique Identifier). A virtual device appears as a real, physical device. You can enumerate up to 128 devices for a virtual device. There is a limit of 4095 on the number of virtual devices that you can create in a single VSAN.

Note For Cisco MDS SAN-OS Release 3.1(2) and later, SDV has been tested to work with up to 1024 virtual devices per VSAN.

Figure 20-5 shows a configuration that includes a new virtual device, vt1.

Figure 20-5 Creating a Virtual Device

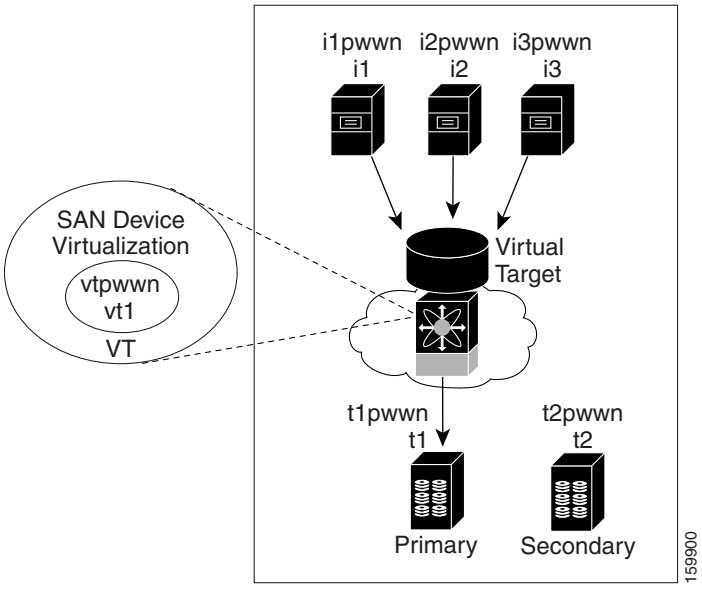

To configure a virtual device and commit it to the fabric configuration, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# sdv enable`	Enables the SDV feature. **Note** If you do not have the Cisco MDS 9000 Family Enterprise package license (ENTERPRISE_PKG) installed, this command will fail.
Step 3	`switch(config)# sdv virtual-device name vdev1 vsan 2`	Configures a virtual device alias name (vdev1). Enters SDV manager configuration submode.
Step 4	`switch(config-sdv-virt-dev)# pwwn 21:00:00:04:cf:cf:45:40 primary` `switch(config-sdv-virt-dev)# pwwn 21:00:00:04:cf:cf:38:d6`	Maps primary virtual device to the pWWN of real devices. Maps secondary virtual device to pWWN of a real device.
Step 5	`switch(config-sdv-virt-dev)# exit`	Exits SDV manager configuration submode.
Step 6	`switch(config)# sdv commit vsan 2`	Commits the VSAN configuration to the fabric.
Step 7	`switch(config)# exit` `switch# show device-alias database`	Verifies that the virtualized device is registered in the device alias database.
Step 8	`switch# show fcns database vsan vsan 2`	If the primary device is online, verifies that the virtual device appears in the name server database and has the correct FC4 type (Fibre Channel layer 4).

Configuring a Zone for a Virtual Device

After configuring a virtual device, you must create a zone that includes all the other real devices and the virtual device as members, and add this zone to a zone set, which you can activate. You can add the virtual device to the zone using the configured name and member type as the device alias.

Note This configuration process does not support interoperability. If you are working in interop-VSANs, we recommend that you configure the zone directly using the system-assigned pWWN of the virtual device.

Figure 20-6 shows a virtual device-name device alias (vt1) zoned with the real devices activated; the primary device is online.

Figure 20-6 Zoning the Virtual Device with Real Devices

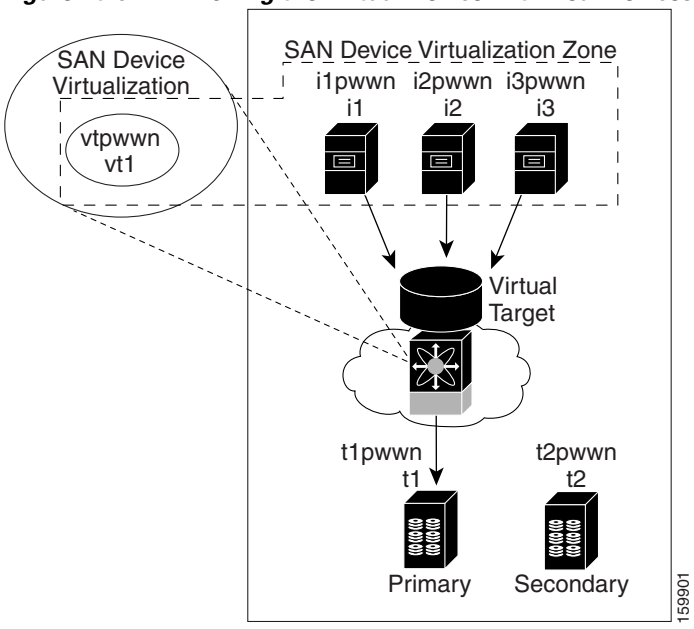

To add the virtual device to a zone as a zone member, follow these steps:

Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zoneset name zs1 vsan 1`	Groups zones under one zone set named zs1 in the VSAN, vsan 1. Enters zone set submode.
Step 3	`switch(config-zoneset)# zone name zone1`	Creates the zone zone1.

Step 4	`switch(config-zoneset-zone)# member device-alias vdev1`
	or
	`switch (config-zoneset-zone)# member pwwn 50:00:53:00:01:fa:70:09`
		Adds a virtual device (vdev1) as device-alias type zone member to the zone.
		Alternatively, adds a virtual device as pWWN type zone member to the zone.

Step 5	`switch(config-zoneset-zone)# member pwwn 21:22:5:d8:15:11:8:8`	Adds the real device to the zone.
Step 6	`switch(config-zoneset-zone)# end`	Exits all modes.

Step 7	`switch# show zoneset`	Displays the zone sets configured for the VSAN, vsan 1, in step 4, as well as the primary and secondary pWWNs.
Step 8	`switch# config t`	Enters configuration mode.
Step 9	`switch(config)# zoneset activate name zs1 vsan 1`	Activates the new zone, zs1 for the VSAN, vsan1.
Step 10	`switch(config)# end`	Exits all modes.
Step 11	`switch# show zoneset active vsan 1`	Displays the active zone for the VSAN, vsan1. Confirms the zone set and zone names.

Caution Set the device alias mode to **enhanced** when using SDV (because the pWWN of a virtual device could change).

For example, SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenable SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. Hence, you would have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or when the pWWN changes.

Be sure you understand how device alias modes work before enabling them. Refer to Chapter 24, "Distributing Device Alias Services" for details and requirements about device alias modes.

Configuring a Virtual Device with a Static FC ID

Typically, the FC ID for the virtual device is assigned by the system. There is a virtual domain reserved and advertised for SDV devices and it is used when allocating the FC ID. The FC ID is registered with the name server and has the same FC4 properties as the primary device it represents.

When a fabric containing a virtual device configuration reboots, the virtual device's domain or FC ID may change; there is no guarantee that the virtual device FC ID will remain the same because it is not a part of the configuration. You can define the FC ID for a virtual device to be static. Configuring a device to have a static FC ID ensures that the same FC ID is used for the virtual device configuration across SAN-OS reboots, and it also provides the following benefits:

- Makes it easier for SAN administrators to plan and design the SAN (for example, administrators can assign virtual domains and FC IDs to be used by virtual devices).
- Improves monitoring and troubleshooting because the same FC ID is assigned to a virtual device on every instance.

Note This procedure is optional, but we recommend it. You can also configure a static domain ID. Only the domain part of the FC ID is static; other values are system-assigned.

If you are using SDV to virtualize devices for either AIX or HP-UX platforms, we recommend you create a static FC ID.

To configure a static FC ID when creating a virtual device, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# sdv virtual-device name vdev1 vsan 2`	Configures a virtual device alias name (vdev1) and defines its primary device (vsan 2). Enters SDV manager configuration submode.
Step 3	`switch(config-sdv-virt-dev)# virtual-fcid 0x960001` `switch(config-sdv-virt-dev)# exit`	Assigns the FC ID 0x960001 to the virtual device. Exits SDV manager configuration submode.
Step 4	`switch(config)# sdv commit vsan 2`	Commits the VSAN configuration to the fabric.

Linking a Virtual Device with a Physical Device

After creating a virtual device and configuring it as part of a zone, you can define the primary device for it using the **link** command, which is also used to fail over to the secondary device.

Note When a link operation fails over to the secondary device, the virtual device is taken offline and then brought online.

Configuring LUN Zone Members for SDV Devices

You can configure LUN zone members for SDV devices. Following are the types of SDV LUN zone configurations for existing real devices and configured SDV devices:

- Real initiator (I1)
- Real target (T1) supporting LUNs from 0 to 12
- Virtual target (VT1) virtualizing the real target (T1)
- Virtual initiator (VI1) virtualizing real initiator (I1)

Real Initiator and SDV Virtual Target with LUN

In Example 20-1 a real initiator is zoned with an SDV virtual target (including the LUN).

Example 20-1 Real Initiator and SDV Virtual Target with LUN

```
zoneset name zs1 vsan 2
    zone name z1 vsan 2
        member device-alias I1
        member device-alias VT1 lun 0
        member device-alias VT2 lun 1
```

SDV Virtual Initiator and Real Target with LUN

In Example 20-2 an SDV virtual initiator is zoned with a real target (including the LUN).

Example 20-2 SDV Virtual Initiator and Real Target with LUN

```
zoneset name zs2 vsan 1
    zone name z2 vsan 1
        member device-alias VI1
        member device-alias T1 lun 0
        member device-alias T2 lun 1
```

SDV Virtual Initiator and SDV Virtual Target with LUN.

In Example 20-3 an SDV virtual initiator is zoned with an SDV virtual target (including the LUN).

Example 20-3

```
zoneset name zs3 vsan 1
    zone name z3 vsan 1
        member device-alias VI1
        member device-alias VT1 lun 0
        member device-alias VT2 lun 1
```

Resolving Fabric Merge Conflicts

Whenever two fabrics merge SDV merges its database. A merge conflict can occur when there is a run-time information conflict or configuration mismatch. Run-time conflicts can occur do to:

- Identical pWWNs being assigned to different virtual devices
- The same virtual devices are assigned different pWWNs.
- The virtual device and virtual FC ID are mismatched.

A *blank commit* is a commit operation that does not contain configuration changes, and enforces the SDV configuration of the committing switch fabric-wide. A blank commit operation resolves merge conflicts by pushing the configuration from the committing switch throughout the fabric, thereby reinitializing the conflicting virtual devices. Exercise caution while performing this operation, as it can easily take some virtual devices offline.

Merge failures resulting from a pWWN conflict can cause a failure with the device alias as well. A blank commit operation on a merge-failed VSAN within SDV should resolve the merge failure in the device alias.

You can avoid merge conflicts due to configuration mismatch by ensuring that:

- The pWWN and device alias entries for a virtual device are identical (in terms of primary and secondary).
- There are no virtual device name conflicts across VSANs in fabrics.

SDV Requirements and Guidelines

Be aware of the following requirements and guidelines as you plan and configure SDV:

- SDV should be enabled on switches where devices that are part of SDV zones are connected.
- SDV does not work for devices connected to non-MDS switches.
- Broadcast zoning is not supported for a zone with a virtual device.

- IVR and SDV cannot be used for the same device. In other words, a SDV-virtualized device cannot be part of a IVR zone or zoneset.
- Virtual device names should be unique across VSANs because they are registered with the device alias server, which is unaware of VSANs. For example, if you have enabled SDV and have registered a name, vt1 in both VSAN 1 and VSAN 2, then the device alias server cannot store both entries because they have the same name.
- You cannot specify the same primary device for different virtual devices.
- SDV does not work with soft zoning (*Soft zoning* means that 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), nor does it work with the **zone default-zone permit vsan** operation (which would otherwise permit or deny traffic to members in the default zone).
- If devices are not already zoned with the initiators, then you can configure SDV virtual device zones with no negative impact. If they are already zoned, then zoning changes are required.
- The real device-virtual device zone cannot coexist with the real device-real device zone. If the real devices are not already zoned together, then you can configure the real device-virtual device zone with no negative impact. If these devices are already zoned, then adding the real device-virtual device zone may cause the zone activation to fail. If this occurs, then you must delete one of the zones before activation.

  For example, a user attempts to create a configuration with zone A, which consists of I, the initiator, and T, the target (I,T), and zone B, which consists of a virtual initiator, VI, and real target, T (zone VI, T). Such a configuration would fail. Likewise, an attempt to configure zone C, which consists of an initiator, I, and target T, with zone D, which consists of an initiator, I, and virtual target, VT (zone I, VT), would also fail.

Caution There must be at least one SDV-enabled switch that is *not* a Cisco MDS 9124 Switch between the server and the device that are being virtualized. In other words, SDV does not work when initiators and primary devices are connected to the same Cisco MDS 9124 Switch.

Discarding Changes

At any time, you can discard the uncommitted changes to the running configuration and release the fabric lock (prior to issuing the **sdv commit** command). If you discard the pending changes, the configuration remains unaffected and the lock is released.

To discard uncommitted SDV configuration changes and release the lock follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# sdv abort vsan 10`	Discards the pending configuration changes.

Clearing SDV Changes

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

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

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

```
switch# clear sdv session vsan 2
```

Guidelines for Downgrading SDV

Prior to SAN-OS Release 3.1(3), SDV did not support virtual initiators and LUN zoning. Consequently, in SAN-OS Releases 3.1(3) and later, if virtual initiators are configured or SDV devices are configured as LUN-based members of a zone, a configuration check will indicate that downgrading to SAN-OS Release 3.1(2) may be disruptive and is therefore not recommended.

Downgrading With Virtual Initiators Configured

If SDV virtual initiators are configured, you will be unable to downgrade to SAN-OS release 3.1(2)..

This incompatibility is a *loose* one–it will only warn users before a downgrade. It is recommended that you remove the virtual initiator configuration or shut down the initiator port so that there are no inconsistencies in the downgraded version.

Note We also recommend that users trigger a manual discovery on all the switches before configuring the virtual initiators; in fact, you can trigger the discovery before proceeding to the downgrade by entering the following commands.

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

Downgrading with SDV LUN Zoning Configured

Following are the downgrade scenarios when SDV LUN zoning is configured:

- Real initiator and SDV virtual target with LUN
- SDV virtual initiator and real target with LUN
- SDV virtual initiator and SDV virtual target with LUN

In each of these cases, a configuration check is registered to prevent users from downgrading to SAN-OS Release 3.1(2). This incompatibility is a *strict* one (it will be disruptive if you proceed with the downgrade).

To avoid the configuration check, delete all the LUN zone members from SDV zones and then activate the zone set before the downgrade.

SDV Configuration Example

The following example shows all tasks (required and optional) associated with configuring SDV.

Step 1 Enter configuration mode.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
```

Step 2 Enable SDV.

```
switch(config)# sdv enable
```

Step 3 Locate the pWWNs of the devices in the VSAN (vsan 2). Record the pWWNs, as you will need them in the next step when you create a virtual device.

```
switch(config)# do show fcns database vsan 2
VSAN 2:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x9f0201    NL    21:00:00:04:cf:cf:45:40 (Seagate)       scsi-fcp
0x9f0423    NL    21:00:00:04:cf:cf:38:d6 (Seagate)       scsi-fcp

Total number of entries = 2
```

Step 4 Create a virtual device (vdev1) for the VSAN and specify both the primary and secondary pWWNs.

```
switch(config)# sdv virtual-device name vdev1 vsan 2
switch(config-sdv-virt-dev)# pwwn 21:00:00:04:cf:cf:45:40 primary
switch(config-sdv-virt-dev)# pwwn 21:00:00:04:cf:cf:38:d6
```

Step 5 Create a static FC ID for the target device.

```
switch(config-sdv-virt-dev)# virtual-fcid 0x960001
switch(config-sdv-virt-dev)# exit
```

Step 6 Commit the new virtual device configuration to the fabric.

```
switch(config)# sdv commit vsan 2
switch(config)# exit
```

Step 7 Enter the **show sdv database** command, which displays the primary and secondary virtual devices created when you configured the virtual device in the VSAN (vsan 2). Ensure that the virtual device name, pWWNs, and FC ID are correct.

```
switch# show sdv database vsan 2
virtual-device name vdev1 vsan 2
[ WWN:50:00:53:00:00:d2:e0:01 FCID:0x960001 Real-FCID:0x9f0201 ]
  virtual-fcid 0x960001
  pwwn 21:00:00:04:cf:cf:45:40 primary
  pwwn 21:00:00:04:cf:cf:38:d6
```

Step 8 Enter the **show device-alias database** command, which displays the contents of the device alias database. Ensure that the new virtual device name appears and that the name is correct.

```
switch# show device-alias database

device-alias name vdev1 pwwn 50:00:53:00:01:c9:70:01

Total number of entries = 1
```

Step 9 Create a zone member under one zone set named zs1 in the VSAN (vsan 2) and add the virtual target device to the new zone using the device alias member type. Before activating the new zone, display the zone-set information to ensure that the zone-set name, zone name, and pWWNs are correct.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# zoneset name zs1 vsan 2
switch(config-zoneset)# zone name zzz1
switch(config-zoneset-zone)# member device-alias vdev1
switch(config-zoneset-zone)# member pwwn 21:00:03:04:55:cf:d6:40
switch(config-zoneset-zone)# end
switch# show zoneset
zoneset name zs1 vsan 2
  zone name zzz1 vsan 2
    pwwn 50:00:53:00:01:c9:70:01 [vdev1]
    pwwn 21:00:03:04:55:cf:d6:40
```

Step 10 Activate the new zone configuration.

```
switch(config)# zoneset activate name zs1 vsan 2
Zoneset activation initiated. check zone status
switch(config)# exit
```

Step 11 Display the active zone-set to ensure the data in the new zone configuration is correct. Also confirm that the pWWNs are correct.

```
switch# show zoneset active vsan 2
zoneset name zs1 vsan 2
  zone name zzz1 vsan 2
  * fcid 0x211324 [pwwn 50:00:53:00:01:c9:70:01] [vdev1]
    pwwn 21:00:03:04:55:cf:d6:40

switch# show sdv database vsan 2
virtual-device name vdev1 vsan 2
[ WWN:50:00:53:00:00:d2:e0:01 FCID:0x960001 Real-FCID:0x9f0201 ]
  virtual-fcid 0x960001
  pwwn 21:00:00:04:cf:cf:45:40 primary
  pwwn 21:00:00:04:cf:cf:38:d6
```

Step 12 Enter the SDV manager configuration submode and display the pWWNs for the virtual devices in vsan 2. After making a note of the pWWNs, migrate the primary virtual device to the secondary device, link this new primary device to a physical device, and then commit this configuration to the fabric. Confirm that the secondary virtual device is now the primary.

```
switch# conf t

switch(config-sdv-virt-dev)# link pwwn 21:00:00:04:cf:cf:38:d6
switch(config-sdv-virt-dev)# exit
switch(config)# sdv commit vsan 2
switch(config)# do show sdv database vsan 2
virtual-device name vdev1 vsan 2
[ WWN:50:00:53:00:00:d2:e0:01 FCID:0x960001 Real-FCID:0x9f0423 ]
  virtual-fcid 0x960001
  pwwn 21:00:00:04:cf:cf:45:40
  pwwn 21:00:00:04:cf:cf:38:d6 primary
```

```
virtual-device name vdev2 vsan 2
[ WWN:50:00:53:00:00:0b:50:01 ]
```

Displaying SDV Information

To display the results of the last commit to the SDV database:

```
switch# show sdv session status vsan  1
Session Parameters for VSAN: 1
------------------------------
Last Action Time Stamp     : Fri Feb  2 10:17:20 2007
Last Action                : Commit
Last Action Result         : Success
Last Action Failure Reason : none
```

To display the results of the last CFS SDV fabric merge for a VSAN:

```
switch# show sdv merge status vsan  1
Merge Status for VSAN    : 1
------------------------------
Last Merge Time Stamp    : None
Last Merge State         : None
Last Merge Result        : SUCCESS
Last Merge Failure Reason: None [cfs_status: 0]
```

To display details about the SDV database:

```
switch# show sdv database vsan 2
virtual-device name vdev1 vsan 2
[ WWN:50:00:53:00:00:d2:e0:01 FCID:0x960001 Real-FCID:0x9f0201 ]
  virtual-fcid 0x960001
  pwwn 21:00:00:04:cf:cf:45:40 primary
  pwwn 21:00:00:04:cf:cf:38:d6
```

To display statistics about SDV for a VSAN:

```
switch# show sdv statistics vsan 1
VSAN ELS-CMD    Requests  Accepts  Rejects  Drops
-----------------------------------------------
1    ELS_PLOGI   54        54         0         0
1    ELS_RRQ      2         2         0         0
1    ELS_PRLI    54        54         0         0
```

Default Settings

Table 20-1 lists the default settings for SDV parameters.

Table 20-1 Default SDV Configuration Parameters

Parameters	Default
enable	disabled

CHAPTER 21

Creating Dynamic VSANs

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

You can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as Dynamic Port VSAN Membership (DPVM). DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved. To assign VSANs statically, see Chapter 19, "Configuring and Managing VSANs."

This chapter includes the following sections:

DPVM

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco SAN-OS software 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 consisting of multiple devices. You can assign any one of these identifiers or any combination of these identifiers to configure DPVM mapping. If you assign a combination, then preference is given to the pWWN.

DPVM uses the Cisco Fabric Services (CFS) infrastructure to allow efficient database management and distribution. DPVM uses the application driven, coordinated distribution mode and the fabric-wide distribution scope (see Chapter 6, "Using the CFS Infrastructure").

Note DPVM does not cause any changes to device addressing. DPVM only pertains to the VSAN membership of the device, ensuring that the host gets same VSAN membership on any port on the switch. For example, if a port on the switch has a hardware failure, you can move the host connection to another port on the switch and not need to update the VSAN membership manually.

Note DPVM is not supported on FL ports. DPVM is supported only on F ports.

This section describes DPVM and includes the following topics:

About DPVM Configuration

To use the DPVM feature as designed, be sure to verify the following requirements:

- The interface through which the dynamic device connects to the Cisco MDS 9000 Family switch must be configured as an F port.
- The static port VSAN of the F port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).

Note The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

Enabling DPVM

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

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

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm enable`	Enables DPVM on that switch.
	`switch(config)# no dpvm enable`	Disables (default) DPVM on that switch.

About DPVM Databases

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

The DPVM feature uses three databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database when distribution is disabled.
- Active database—The database currently enforced by the fabric.
- Pending database—All configuration changes are stored in the DPVM pending database when distribution is enabled (see the "DPVM Database Distribution" section on page 21-5).

Changes to the DPVM config database are not reflected in the active DPVM database until you activate the DPVM config database. Changes to the DPVM pending database are not reflected in the config/active DPVM database until you commit the DPVM pending database. This database structure allows you to create multiple entries, review changes, and let the DPVM config and pending databases take effect.

Configuring DPVM Config and Pending Databases

To create and populate the DPVM config and pending databases, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm database` `switch(config-dpvm-db)#`	Creates the DPVM config database.
	`switch(config)# no dpvm database`	Deletes the DPVM config database.
Step 3	`switch(config-dpvm-db)# pwwn 12:33:56:78:90:12:34:56 vsan 100`	Maps the specified device pWWN to VSAN 100.
	`switch(config-dpvm-db)# no pwwn 12:33:56:78:90:12:34:56 vsan 101`	Removes the specified device pWWN mapping from the DPVM config database.
Step 4	`switch(config-dpvm-db)# nwwn 14:21:30:12:63:39:72:81 vsan 101`	Maps the specified device nWWN to VSAN 101.
	`switch(config-dpvm-db)# no nwwn 14:21:30:12:63:39:72:80 vsan 101`	Removes the specified device nWWN mapping from the DPVM config database.

Activating DPVM Config Databases

When you explicitly activate the DPVM config database, the DPVM config database becomes the active DPVM database. Activation may 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 issuing the **no dpvm activate** command.

To activate the DPVM config database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm activate`	Activates the DPVM config database.
	`switch(config)# no dpvm activate`	Deactivates the currently active DPVM database.
	`switch(config)# dpvm activate force`	Forcefully activates the DPVM config database to override conflicting entries.

About Autolearned Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. The autolearn feature can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the active DPVM database. The active DPVM database should already be available to enable autolearn.

You can delete any learned entry from the active DPVM database when you enable autolearn. These entries only become permanent in the active DPVM database when you disable autolearn.

Note Autolearning is only supported for devices connected to F ports. Devices connected to FL ports are not entered into the DPVM database because DPVM is not supported on FL ports.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the active DPVM database.
- If the same device logs multiple times into the switch through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process—enabling autolearning followed by disabling autolearning.When the **auto-learn** option is enabled, the following applies:
  - Learning currently logged-in devices—occurs from the time learning is enabled.
  - Learning new device logins— occurs as and when new devices log in to the switch.

Enabling Autolearning

To enable autolearning, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm auto-learn`	Enables learning on this switch.
	`switch(config)# no dpvm auto-learn`	Disables (default) learning on this switch.

Clearing Learned Entries

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

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

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

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

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

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

DPVM Database Distribution

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement (see Chapter 6, "Using the CFS Infrastructure").

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

About DPVM Database Distribution

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.

If fabric distribution is enabled, all changes to the configuration database are stored in the DPVM pending database. These changes include the following tasks:

- Adding, deleting, or modifying database entries.
- Activating, deactivating, or deleting the configuration database.
- Enabling or disabling autolearning.

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.

Tip You can view the contents of the DPVM pending database by issuing the **show dpvm pending** command.

Disabling DPVM Database Distribution

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no dpvm distribute`	Disables DPVM distribution to the neighboring switches.
	`switch(config)# dpvm distribute`	Enables (default) DPVM distribution to the neighboring switches.

About Locking the Fabric

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

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the DPVM pending database. Modifications from this point on are made to the DPVM pending database. The DPVM pending database remains in effect until you commit the modifications to the DPVM pending database or discard (abort) the changes to the DPVM pending database.

Locking the Fabric

To lock the fabric and apply changes to the DPVM pending database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm database` `switch(config-dpvm-db)#`	Accesses the DPVM config database.
Step 3	`switch(config-dpvm-db)# pwwn 11:22:33:44:55:66:77:88 vsan 11`	Adds one entry to the DPVM config database.

	Command	Purpose
Step 4	`switch(config-dpvm-db)# exit` `switch(config)#`	Exits to configuration mode.
Step 5	`switch(config)# dpvm activate`	Activates the DPVM config database.

Committing Changes

If you commit the changes made to the configuration, the configuration in the DPVM pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the DPVM pending database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm commit`	Commits the database entries that are currently in the DPVM pending database.

Discarding Changes

If you discard (abort) the changes made to the DPVM pending database, the configurations remain unaffected and the lock is released.

To discard the DPVM pending database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm abort`	Discards the database entries that are currently in the DPVM pending database.

Clearing a Locked Session

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.

Tip The DPVM 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 dpvm session** command in EXEC mode.

```
switch# clear dpvm session
```

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active DPVM database. See the "CFS Merge Support" section on page 6-8 for detailed concepts.

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

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

Caution If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

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

About Copying DPVM Databases

The following circumstances may require the active DPVM database to be copied to the DPVM config database:

- If the learned entries are only added to the active DPVM database.
- If the DPVM config database or entries in the DPVM config database are accidently deleted.

Note If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

Copying DPVM Databases

To copy the currently active DPVM database to the DPVM config database, use the **dpvm database copy** command.

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

Comparing Database Differences

You can compare the DPVM databases as follows:

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

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

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

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

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

To add pending database entries to the DPVM config database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# dpvm distribute`	Enables CFS distribution.

	Command	Purpose
Step 3	`switch(config)# dpvm database`	Accesses the DPVM config database.
Step 4	`switch(config-dpvm-db)# pwwn 44:22:33:44:55:66:77:88 vsan 55` `switch(config-dpvm-db)# pwwn 55:22:33:44:55:66:77:88 vsan 55`	Adds two entries to the DPVM config database.

Displaying DPVM Configurations

Use the **show dpvm** command to display information about WWNs configured on a per VSAN basis (see Examples 21-1 to 21-6).

Example 21-1 Displays the DPVM Configuration Status

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

Example 21-2 Displays the DPVM Current Dynamic Ports for the Specified VSAN

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

Example 21-3 Displays the DPVM Config Database

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

Example 21-4 Displays the DPVM Database

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

Example 21-5 Displays DPVM Config Database

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

Example 21-6 Compares Pending Database with the DPVM Config Database

```
switch# show dpvm pending-diff
```

```
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----------------------------------------------------------------------
+  pwwn  55:22:33:44:55:66:77:88  vsan 55
-  pwwn  11:22:33:44:55:66:77:88  vsan 11
*  pwwn  44:22:33:44:55:66:77:88  vsan 44
```

Sample DPVM Configuration

To configure a basic DPVM scenario, follow these steps:

Step 1 Enable DPVM and enable DPVM distribution.

```
switch1# config
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# dpvm enable
switch1(config)# end
switch1# show dpvm database
switch1# show dpvm database active
switch1# show dpvm status

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

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

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

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

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

```
switch1# config
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end
switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4(*)
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5(*)
[Total 2 entries]
* is auto-learnt entry
switch1# show dpvm ports
------------------------------------------------------------
Interface   Vsan        Device pWWN       Device nWWN
------------------------------------------------------------
fc1/24       4   21:00:00:e0:8b:0e:74:8a  20:00:00:e0:8b:0e:74:8a
fc1/27       5   21:01:00:e0:8b:2e:87:8a  20:01:00:e0:8b:2e:87:8a
switch1# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID            PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/24     4     0xe70100  21:00:00:e0:8b:0e:74:8a  20:00:00:e0:8b:0e:74:8a
fc1/27     5     0xe80100  21:01:00:e0:8b:2e:87:8a  20:01:00:e0:8b:2e:87:8a

Total number of flogi = 2.
```

```
switch195# show dpvm status
DB is activated successfully, auto-learn is on
```

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

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

Step 4 Access switch9 and issue the following commands.

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

Step 5 Access switch3 and issue the following commands.

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

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

```
switch1# config
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# no dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end
switch1# show dpvm status
DB is activated successfully, auto-learn is off
switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
[Total 6 entries]
* is auto-learnt entry
switch1# show dpvm status
DB is activated successfully, auto-learn is off
```

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

Step 7 Access switch9 and issue the following commands.

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

Step 8 Access switch3 and issue the following commands.

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

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

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

Default Settings

Table 21-1 lists the default settings for DPVM parameters.

Table 21-1 Default DPVM Parameters

Parameters	Default
DPVM	Disabled.
DPVM distribution	Enabled.
Autolearning	Disabled.

CHAPTER 22

Configuring Inter-VSAN Routing

This chapter explains the Inter-VSAN routing (IVR) feature and provides details on sharing resources across VSANs using IVR management interfaces provided in the switch.

This chapter includes the following sections:

Inter-VSAN Routing

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 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.

This section includes the following topics:

About IVR

Note IVR is not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. Fibre Channel control traffic does not flow between VSANs, nor can initiators access any resource across VSANs other than the designated ones. Valuable resources such as tape libraries are easily shared across VSANs without compromise.

IVR is in compliance with Fibre Channel standards and incorporates third-party switches, however, IVR-enabled VSANs may have to be configured in one of the interop modes.

IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions (see Figure 22-1).

Note See the "Example Configurations" section on page 22-39 for procedures to configure the sample scenario shown in Figure 22-1.

Figure 22-1 Traffic Continuity Using IVR and FCIP

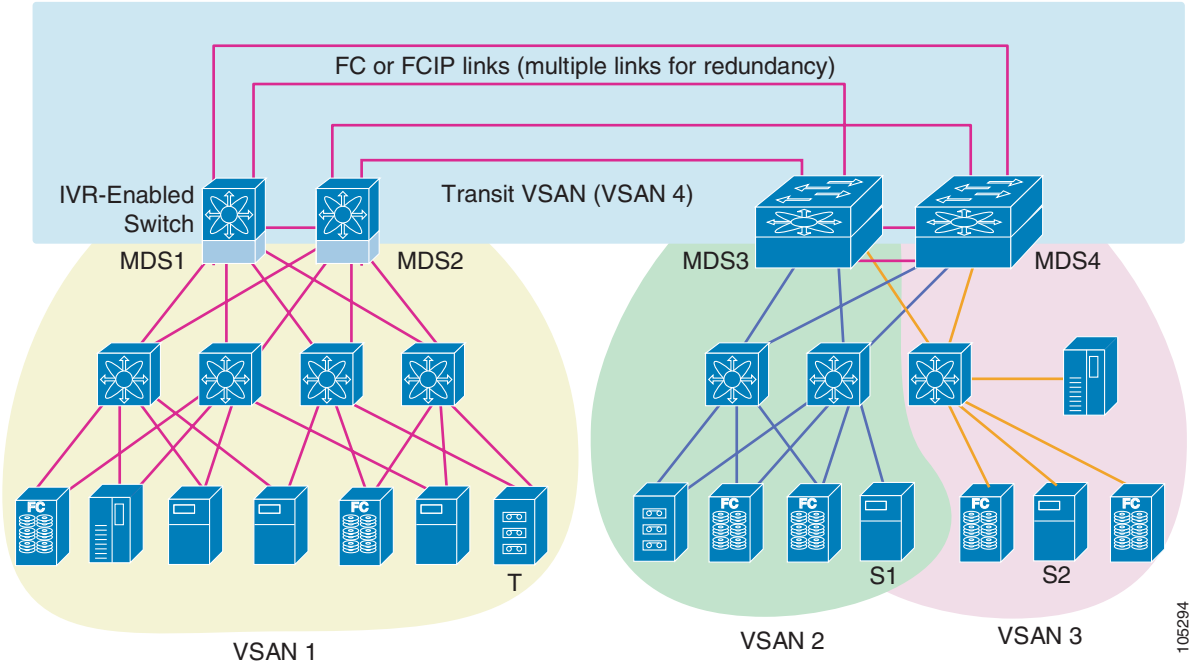

Note OX ID based load balancing of IVR traffic from IVR- enabled switches is not supported on Generation 1 switching modules. OX ID based load balancing of IVR traffic from a non-IVR MDS switch should work. Generation 2 switching modules support OX ID based load balancing of IVR traffic from IVR-enabled switches.

IVR Features

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.
- Shares valuable resources (like tape libraries) across VSANs without compromise.
- 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 have to be configured in one of the interop modes.

IVR Terminology

The following IVR-related terms are used in this chapter.:

- 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)—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. Prior to Cisco SAN-OS Release 3.0(3), you can configure up to 2000 IVR zones and 10,000 IVR zone members on the switches in the network. As of Cisco SAN-OS Release 3.0(3), you can configure up to 8000 IVR zones and 20,000 IVR zone members on the switches in the network.
- Inter-VSAN routing zone sets (IVR zone sets)—One or more IVR zones make up an IVR zone set. You can configure up to 32 IVR zone sets on any switch in the Cisco MDS 9000 Family. Only one IVR zone set can be active at any time.
- 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. In Figure 22-1, VSANs 1, 2, and 3 are edge VSANs.

  Note An edge VSAN for one IVR path can be a transit VSAN for another IVR path.

- 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. In Figure 22-1, VSAN 4 is a transit VSAN.

  Note When the source and destination edge VSANs are adjacent to each other, then a transit VSAN is not required between them.

- Border switch—An IVR-enabled switch that is a member of two or more VSANs. Border switches, such as the IVR-enabled switch between VSAN 1 and VSAN 4 in Figure 22-1, span two or more different color-coded VSANs.
- Edge switch—A switch to which a member of an IVR zone has logged in. Edge switches are unaware of the IVR configurations in the border switches. Edge switches need not 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 enabling 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.

IVR Limits Summary

Table 22-1 summarizes the configuration limits for IVR. See Appendix A, "Configuration Limits for Cisco MDS SAN-OS Release 3.x," for a complete list of Cisco MDS SAN-OS feature configuration limits.

Table 22-1 IVR Configuration Limits

IVR Feature	Maximum Limit
IVR zone members	20,000 IVR zone members per physical fabric as of Cisco SAN-OS Release 3.0(3). 10,000 IVR zone members per physical fabric prior to Cisco SAN-OS Release 3.0(3).
IVR zones	8000 IVR zones per physical fabric as of Cisco SAN-OS Release 3.0(3). 2000 IVR zones per physical fabric prior to Cisco SAN-OS Release 3.0(3).
IVR zone sets	32 IVR zone sets per physical fabric.
IVR service groups	16 service groups per physical fabric.

Fibre Channel Header Modifications

IVR works by virtualizing 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 goes 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.

IVR NAT

Without Network Address Translation (NAT), IVR requires unique domain IDs for all switches in the fabric. You can enable IVR NAT to allow non-unique domain IDs. This feature simplifies the deployment of IVR in an existing fabric where non-unique domain IDs might be present.

To use IVR NAT, it must be enabled in all IVR-enabled switches in the fabric IVR configuration distribution (see the "Distributing the IVR Configuration using CFS" section on page 22-10). By default, IVR NAT and IVR configuration distribution are disabled in all switches in the Cisco MDS 9000 Family.

IVR NAT Requirements and Guidelines

Following are requirements and guidelines for using IVR NAT:

- For IVR NAT to function correctly in the network, all IVR-enabled switches must run Cisco MDS SAN-OS Release 2.1(1a) or later.
- IVR NAT port login (PLOGI) requests 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. However, load balancing of IVR NAT traffic over PortChannel links is supported. The load balancing algorithm for IVR NAT traffic over port-channel with Generation 1 linecards is SRC/DST only. Generation 2 linecards support SRC/DST/OXID based load balancing of IVR NAT traffic across a port-channel.
- You cannot configure IVR NAT and preferred Fibre Channel routes on Generation 1 module interfaces.

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 destinations IDs are part of the payload. 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 described in Table 22-2.

Table 22-2 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

Table 22-2 Extended Link Service Messages Supported by IVR NAT (continued)

Extended Link Service Messages	Link Service Command (LS_COMMAND)	Mnemonic
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

If you have a message that is not recognized by IVR NAT and contains the destination ID in the payload, you cannot use IVR with NAT in your topology. You can still use IVR with unique domain IDs.

IVR VSAN Topology

IVR uses a configured IVR VSAN topology to determine how to route traffic between the initiator and the target across the fabric. You can configure this IVR VSAN topology manually on an IVR-enabled switch and distribute the configuration using CFS in Cisco MDS SAN-OS Release 2.0(1b) or later. Alternately, in Cisco MDS SAN-OS Release 2.1(1a) or later, you can configure IVR topology in auto mode. Prior to Cisco MDS SAN-OS Release 2.0(1b), you need to manually copy the IVR VSAN topology to each switch in the fabric.

Auto mode automatically builds the IVR VSAN topology and maintains the topology database when fabric reconfigurations occur. Auto mode distributes the IVR VSAN topology to IVR-enabled switches using CFS.

Using auto mode, you no longer need to manually update the IVR VSAN topology when reconfigurations occur in your fabric. If a manually configured IVR topology database exists, auto mode initially uses that topology information. This 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 auto IVR topology is turned on it starts with the previously active, if any, manual IVR topology. Auto topology then commences its discovery process, and may come up with new, alternate or better paths. If the traffic is switched to an alternate or better path, there may be temporary traffic disruptions that are normally associated with switching paths.

Note IVR topology in auto mode requires Cisco MDS SAN-OS Release 2.1(1a) or later and enabling CFS for IVR on all switches in the fabric.

Autonomous Fabric ID

The autonomous fabric ID (AFID) distinguishes segmented VSANS (that is, two VSANs that are logically and physically separate but have the same VSAN number). Cisco MDS SAN-OS supports AFIDs from 1 through 64. AFIDs are used in conjunction with auto mode to allow segmented VSANS in the IVR VSAN topology database. You can configure up to 64 AFIDs.

The AFID can be configured individually for each switch and list of VSANs, or the default AFID can be configured for each switch.

Note Two VSANs with the same VSAN number but different AFIDs are counted as two VSANs out of the total 128 VSANs allowed in the fabric.

IVR Service Groups

IVR service groups have the following characteristics:

- 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/AFID combination cannot be a member of more than one service group. CFS merge fails if such a condition exists.
- Total number of AFID/VSAN combinations in all the service groups combined cannot exceed 128. The maximum number of AFID/VSAN combinations in a single service group is 128.
- IVR control traffic is distributed in all the members of all the service groups. IVR data traffic between two end devices belonging to a service group stays within that service group. For example, two members pWWN 1 and pWWN 2 belonging to the same IVR zone but different service groups cannot communicate.
- During a CFS merge, service groups with 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 the reachable SANs. If you do not enable CFS distribution, you must ensure that the service group configuration is same at all the IVR-enabled switches in all the VSANs.
- IVR end devices belonging to an IVR service group are not exported to any AFID/VSAN outside of its service group.
- If at least one service group is defined and an IVR zone member that does not belong to a service group, that IVR zone member is not able to communicate with any other device.
- The default service group ID is zero (0).

Default Service Group

All AFID/VSAN combinations that are part of 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.

Service Group Activation

A configured service group must be activated for it take effect. Like zoneset activation or VSAN topology activation, the activation of a configured service group replaces the currently activate 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.

IVR Interoperability

When using the IVR feature, all border switches in a given fabric must be Cisco MDS switches. However, other switches in the fabric may be non-MDS switches. For example, end devices that are members of the active IVR zone set may be connected to non-MDS switches. Non-MDS switches may also be present in the transit VSAN(s) or in the edge VSANs if one of the **interop** modes is enabled.

See the "Switch Interoperability" section on page 29-11.

IVR Configuration Task List

To configure IVR in a SAN fabric, follow these steps:

Step 1 Determine whether to use IVR Network Address Translation (NAT).

Step 2 If you do not plan to use IVR NAT, verify that unique domain IDs are configured in all switches and VSANs participating in IVR.

Step 3 Enable IVR in the border switches.

Step 4 Configure the service group as required.

Step 5 Configure fabric distribution as required.

Step 6 Configure the IVR topology, either manually or automatically.

Step 7 Create and activate IVR zone sets in *all* of the IVR-enabled border switches, either manually or using fabric distribution.

Step 8 Verify the IVR configuration.

Configuring IVR

This section describe how to configure IVR and contains the following sections:

Enabling IVR

The IVR feature must be enabled in all border switches in the fabric that participate in the IVR. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. You can manually enable IVR on all required switches in the fabric or configure fabric-wide distribution of the IVR configuration ("Distributing the IVR Configuration using CFS" section on page 22-10).

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

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ivr enable`	Enables IVR on the switch.
	`switch(config)# no ivr enable`	Disables (default) IVR on the switch.

Distributing the IVR Configuration using CFS

The IVR feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient configuration management and to provide a single point of configuration for the entire fabric in the VSAN (see Chapter 6, "Using the CFS Infrastructure").

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).
- IVR service groups.
- AFID database.

Note IVR configuration distribution is disabled by default. For the feature to function correctly, you must enable it on all IVR-enabled switches in the network.

This section includes the following topics:

Database Implementation

The IVR feature uses three databases to accept and implement configurations.

- Configured database—The database is manually configured by the user.
- Active database—The database is currently enforced by the fabric.
- Pending database—If you modify the configuration, you need to commit or discard the configured database changes to the pending database. The fabric remains locked during this period. Changes to the pending database are not reflected in the active database until you commit the changes to CFS.

Enabling Configuration Distribution

To enable IVR configuration distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr distribute`	Enables IVR distribution.
	`switch(config)# no ivr distribute`	Disables (default) IVR distribution.

Locking the Fabric

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

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

Committing the Changes

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

To commit IVR configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr commit`	Commits the IVR changes.

Discarding the Changes

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

To discard IVR configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr abort`	Discards the IVR changes and clears the pending configuration database.

Clearing a Locked Session

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

Tip The 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 ivr session** command in EXEC mode.

```
switch# clear ivr session
```

About IVR NAT and Auto Topology

Before configuring an IVR SAN fabric to use IVR NAT and auto-topology, consider the following guidelines:

- Configure IVR only in the relevant switches.
- Enable CFS for IVR on all switches in the fabric.
- Verify that all switches in the fabric are running Cisco MDS SAN-OS Release 2.1(1a) or later.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package if you have Cisco MDS SAN-OS Release2.1(1a) or later and one active IPS card for this feature (see Chapter 3, "Obtaining and Installing Licenses").

Note The IVR over FCIP feature is bundled with the Cisco MDS 9216i Switch and does not require the SAN extension over IP package for the fixed IP ports on the supervisor module.

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.

Note IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

Transit VSAN Guidelines

Consider the following guidelines for transit VSANs:

- 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.
- 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.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 2.1(1a) or later.
- A border switch must be a member of two or more VSANs.

- A border switch that facilities 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.

Service Group Guidelines

If you use service groups with IVR auto topology, you should enable IVR and configure your service groups first, then distribute them with CFS before setting the IVR topology in auto mode.

Configuring IVR Topology Automatic Mode

Note IVR configuration distribution must be enabled before configuring IVR topology automatic mode (see the "Distributing the IVR Configuration using CFS" section on page 22-10). Once IVR topology automatic mode is enabled, you cannot disable IVR configuration distribution.

To configure IVR topology automatic mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr vsan-topology auto`	Configures IVR topology automatic mode.
	`switch(config)# ivr vsan-topology activate`	Disables IVR topology automatic mode and reverts to user-configuration mode.

View automatically discovered IVR topology using the **show ivr vsan-topology** command.

```
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 AUTO
Last activation time: Mon Mar 24 07:19:53 1980
```

Note The asterisk (*) indicates the local switch.

Enabling IVR NAT

To configure IVR NAT, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ivr nat`	Enables IVR NAT on the switch.
	`switch(config)# no ivr nat`	Disables (default) IVR NAT on the switch.

About IVR Service Groups

In a complex network topology, you might have only 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.

CFS distribution of IVR information is restricted within the service group only when IVR VSAN topology is in automatic mode. See the "IVR VSAN Topology" section on page 22-6

Configuring IVR Service Groups

To configure an IVR service group, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr service-group name IVR-SG1` `switch(config-ivr-sg)#`	Configures the IVR service group called IVR-SG1 and enters IVR server group configuration mode.
	`switch(config)# no ivr service-group name IVR-SG1` `Successfully erased service group IVR-SG1`	Deletes the IVR service group.
Step 3	`switch(config-ivr-sg)# autonomous-fabric-id 10 vsan-ranges 1,2,6-10`	Configures AFID 10 for VSANs 1, 2, and 6 through 10.
	`switch(config-ivr-sg)# autonomous-fabric-id 11 vsan-ranges 1`	Configures AFID 11 for VSAN 1.
	`switch(config-ivr-sg)# autonomous-fabric-id 12 vsan-ranges 3-5`	Configures AFID 12 for VSANs 3 through 5.
	`switch(config-ivr-sg)# no autonomous-fabric-id 12 vsan-ranges 3-5`	Removes the association between AFID 12 and VSANs 3 through 5.
	`switch(config-ivr-sg)# exit` `switch(config)#`	Returns to configuration mode.
	`switch(config)# ivr service-group name IVR-SG2` `switch(config-ivr-sg)#`	Configures the IVR service group called IVR-SG2 and enters IVR server group configuration mode.
	`switch(config-ivr-sg)# autonomous-fabric-id 20 vsan-ranges 3-5`	Configures AFID 20 for VSANs 3 through 5.

	Command	Purpose
Step 4	`switch(config-ivr-sg)# exit` `switch(config)#`	Returns to configuration mode.
Step 5	`switch(config)# ivr service-group activate`	Activates the service group configuration and sets the communication policy between switches in the default service group as allow (default).
	`switch(config)# ivr service-group activate default-sg-deny`	Activates the service group configuration and 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.
	`switch(config)# no ivr service-group activate`	Deactivates (default) the service group configuration.
Step 6	`switch(config)# ivr vsan-topology activate`	Activates the VSAN topology.
Step 7	`switch(config)# ivr distribute`	Enables CFS distribution for the IVR configuration.
Step 8	`switch(config)# ivr commit`	Commits the IVR configuration to the fabric.

Copying the Active IVR Service Group Database

You cannot modify the active IVR service group database. However, you can modify the configured IVR service group database. To copy the active IVR service group database to the manually configure service group database, use the following command in EXEC mode:

```
switch# ivr copy active-service-group user-configured-service-group
```

Clearing IVR Service Group Database

You can clear all entries in the IVR service group database using the **clear ivr service-group database** command in EXEC mode. This command only clears the configured database, not the active database.

```
switch# clear ivr service-group database
```

Verifying IVR Service Group Configuration

Use the **show ivr service-group active** command to view the active IVR service group database.

```
switch# show ivr service-group active

IVR ACTIVE Service Group
====================================

SG-ID  SG-NAME    AFID  VSANS
--------------------------------------------------------------------------------
 1     IVR-SG1      10  1-2,6-10
```

```
 1     IVR-SG1      11  1
 2     IVR-SG2      20  3-5

Total:   3 entries in active service group table
```

Use the **show ivr service-group configured** command to view the configured IVR service group database.

```
switch# show ivr service-group configured

IVR CONFIGURED Service Group
====================================

SG-ID  SG-NAME    AFID  VSANS
--------------------------------------------------------------------------------
 1     IVR-SG1      10  1-2,6-10
 1     IVR-SG1      11  1
 2     IVR-SG2      20  3-5

Total:   3 entries in configured service group table
```

About AFIDs

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. IVR supports a maximum of 64 AFIDs.

Note You can only use AFID configuration when the VSAN topology mode is automatic. In user-configured VSAN topology mode, the AFIDs are specified in the VSAN topology configuration itself and a separate AFID configuration is not needed.

Configuring Default AFIDs

To configure the default AFID, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# autonomous-fabric-id database`	Enters AFID database configuration submode.
Step 3	`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.
	`switch(config-afid-db)# no switch-wwn 20:00:00:0c:91:90:3e:80 default-autonomous-fabric-id 5`	Reverts to the default value (1) for the default AFID.

Configuring Individual AFIDs

To configure individual AFIDs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# autonomous-fabric-id database`	Enters AFID database configuration submode.
Step 3	`switch(config-afid-db)# switch-wwn 20:00:00:0c:91:90:3e:80 autonomous-fabric-id 10 vsan-ranges 1,2,5-8`	Configures an AFID and VSAN range for a switch. The valid range for AFIDs is 1 to 64.
	`switch(config-afid-db)# no switch-wwn 20:00:00:0c:91:90:3e:80 autonomous-fabric-id 10 vsan-ranges 2`	Deletes VSAN 2 from AFID 10.

Verifying the AFID Database Configuration

View the contents of the AFID database using the **show autonomous-fabric-id database** command.

```
switch# show autonomous-fabric-id database

SWITCH WWN                      Default-AFID
------------------------------------------------------------
20:00:00:0c:91:90:3e:80           5


Total:   1 entry in default AFID table

SWITCH WWN                      AFID      VSANS
------------------------------------------------------------
20:00:00:0c:91:90:3e:80           10    1,2,5-8


Total:   1 entry in AFID table
```

About IVR Without IVR NAT or Auto Topology

Before configuring an IVR SAN fabric without IVR in NAT mode or IVR topology in auto mode, 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
- Configure IVR only in the relevant border switches.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package and one active IPS card for this feature.

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.

Note IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

Domain ID Guidelines

Domain IDs must be unique across inter-connected VSANs when not using IVR NAT. To ensure unique domain IDs across inter-connected VSANs, consider these guidelines:

- 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.

Note In a configuration involving IVR without NAT, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology must be configured with static domain IDs.

Transit VSAN Guidelines

Before configuring transit VSANS, consider the following guidelines:

- 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.
- 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.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco 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 facilities 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 must be updated before a border switch is added or removed.

Configuring IVR Without NAT

You must create the IVR topology in every IVR-enabled switch in the fabric if you have not configured IVR topology in auto mode. You can have up to 128 VSANs in an IVR topology. 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. See Figure 22-2.

Figure 22-2 Example IVR Topology with Non-Unique VSAN IDs Using AFIDs

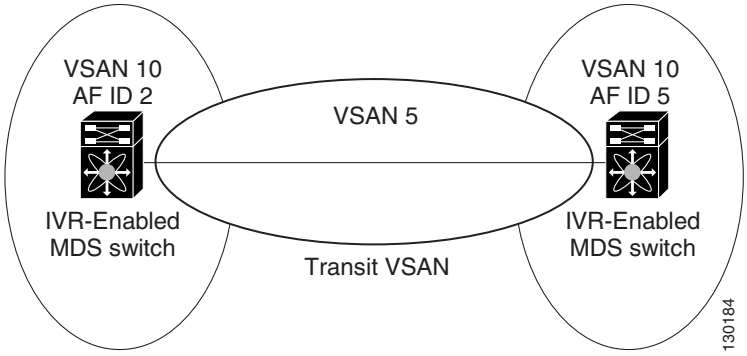

Note If two VSANs in an IVR topology have the same VSAN ID and different AFIDs, they count as two VSANs for the 128-VSAN limit for IVR.

Note The use of a single AFID does not allow for segmented VSANs in an inter-VSAN routing topology.

Caution You can only configure a maximum of 128 IVR-enabled switches and 128 distinct VSANs in an IVR topology (see the "Database Merge Guidelines" section on page 22-37).

Manually Configuring the IVR Topology

Use the sh**ow wwn switch** command to obtain the switch WWNs of the IVR-enabled switches.

To configure a user-defined IVR topology database, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ivr vsan-topology database` `switch(config-ivr-topology-db)#`	Enters the VSAN topology database configuration mode for the IVR feature.
Step 3	`switch(config-ivr-topology-db)# autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:b8 vsan-ranges 1-2,6`	Configures VSANs 1, 2, and 6 to participate in IVR for this switch.
	`switch(config-ivr-topology-db)# autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:c2 vsan-ranges 1-3`	Configures VSANs 1, 2 and 3 to participate in IVR for this switch.
	`switch(config-ivr-topology-db)# no autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:c2 vsan-ranges 1-2`	Removes VSANs 1 and 2 from IVR for this switch.
Step 4	`switch(config-ivr-topology-db)# end` `switch#`	Reverts to EXEC mode.

View your configured IVR topology using the **show ivr vsan-topology** command. 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 *   no      yes  1-2
   1  20:02:00:44:22:00:4a:05     no      yes  1-2,6
   1  20:02:00:44:22:00:4a:07     no      yes  2-5

Total:   3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE
```

Note If CFS is not enabled, you must repeat this configuration in all IVR-enabled switches. See the "Database Merge Guidelines" section on page 22-37.

Tip Transit VSANs are deduced based on your configuration. The IVR feature does not have an explicit transit-VSAN configuration.

Activating a Manually Configured IVR Topology

After manually configuring the IVR topology database, you must activate it.

Caution Active IVR topologies cannot be deactivated. You can only switch to IVR topology automatic mode.

To activate the manually configured IVR topology database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr vsan-topology activate`	Activates the configured IVR topology.

View your active IVR topology using the **show ivr vsan-topology** command.

```
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 1980
```

Note The asterisk (*) indicates the local switch.

Adding an IVR-Enabled Switch to an Existing IVR Topology

Before adding an IVR-enabled switch to an existing fabric with manual IVR topology and CFS distribution enabled (see the "Distributing the IVR Configuration using CFS" section on page 22-10), you must add an entry to the IVR topology for the new switch and activate the new IVR topology.

To add the IVR-enabled switch to the existing IVR topology on the IVR-enabled switch where you update the IVR configuration, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 1	`mds(config)# ivr vsan-topology database` `mds(config-ivr-topology-db)#`	Enters IVR VSAN topology database configuration submode.
Step 2	`mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:05:40:01:1b:c2 vsan-ranges 1,4`	Adds the new IVR-enabled switch to the topology.
Step 3	`switch(config-ivr-topology-db)# exit` `switch(config)#`	Returns to configuration mode.
Step 4	`switch(config)# ivr vsan-topology activate`	Activates the IVR VSAN topology.
Step 5	`switch(config)# ivr commit`	Commits the IVR configuration change to the fabric.
Step 6	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 7	`switch# copy running-config startup-config`	Saves the running configuration.

After adding the switch to the IVR topology, you then enable IVR and CFS for the IVR application on the new switch (see the"Enabling IVR" section on page 22-9 and the "Distributing the IVR Configuration using CFS" section on page 22-10).

Copying the Active IVR Topology

You cannot edit the active IVR topology. However, you can edit the manually configured topology. To copy the active IVR topology database to the manually configure topology, use the following command in EXEC mode:

```
switch# ivr copy active-topology user-configured-topology
```

Clearing the Configured IVR Topology Database

You can only clear manually created IVR VSAN topology entries from the configured database.

To clear the manually configured IVR VSAN topology database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no ivr vsan-topology database`	Clears the previously created IVR topology.

Verifying the IVR Topology

You can verify the IVR topology by using the **show ivr vsan-topology** command. See Example 22-1 to Example 22-3.

Example 22-1 Displays the Configured IVR VSAN Topology

```
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:   5 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Sat Mar 22 21:46:15 1980
```

Note The asterisk (*) indicates the local switch.

Example 22-2 Displays the Active IVR VSAN Topology

```
switch# show ivr vsan-topology active
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:   5 entries in active IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Sat Mar 22 21:46:15
```

Example 22-3 Displays the Configured IVR VSAN Topology

```
switch# show ivr vsan-topology configured
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:   5 entries in configured IVR VSAN-Topology
```

Migrating from IVR Auto Topology Mode to Manual Mode

If you want to migrate the active IVR VSAN topology database from automatic mode to user-configured mode, first copy the active IVR VSAN topology database to the user-configured IVR VSAN topology database before switching modes.

To migrate from automatic mode to manual mode, follow these steps:

	Command	Purpose
Step 1	`switch# ivr copy auto-topology user-configured-topology`	Copies the automatic IVR topology database to the user-configured IVR topology.
Step 2	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 3	`switch(config)# ivr vsan-topology active`	Disabled automatic mode for the IVR topology database and enables user-configuration mode.

About IVR 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 Cisco SN5428) 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(s) 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 domain list for that VSAN.

Tip Be sure to add IVR virtual domains if Cisco SN5428 or MDS 9020 switches exist in the VSAN.

When you enable the IVR virtual domains, links may fail to come up due to overlapping virtual domain identifiers. If so, temporarily withdraw the overlapping virtual domain from that VSAN.

Note Withdrawing an overlapping virtual domain from an IVR VSAN disrupts IVR traffic to and from that domain.

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.

Configuring IVR Virtual Domains

To configure an IVR virtual domain in a specified VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr virtual-fcdomain-add vsan-ranges 1`	Adds the IVR virtual domains in VSAN 1.
	`switch(config)# no ivr virtual-fcdomain-add vsan-ranges 1`	Reverts to the factory default of not adding IVR virtual domains and removes the currently active virtual domains for that VSAN from the fcdomain manger list

Verifying the IVR Virtual Domain Configuration

View the status of the IVR virtual domain configuration using the **show ivr virtual-fcdomain-add-status** command.

```
switch# show ivr virtual-fcdomain-add-status
IVR virtual domains are added to fcdomain list in VSANS: 1
(As well as to VSANs in interoperability mode 2 or 3)
```

Clearing the IVR fcdomain Database

You might want to clear the IVR fcdomain database. You can do this using the following command:

```
switch# clear ivr fcdomain database
```

About Persistent FC IDs for IVR

You can configure persistent FC IDs for IVR. FC ID persistence across reboot improves IVR management by providing the following features:

- Allows you to control and assign a specific virtual domain to use for a native VSAN.

- Allows you to control and assign a specific virtual FC ID to use for a device.

The benefits of persistent FC IDs for IVR are as follows:

- Host devices always see the same FC ID for targets.
- It helps you plan your SAN layout better by assigning virtual domains for IVR to use.
- It 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.

You can configure two types of database entries for persistent IVR FC IDs:

- Virtual domain entries—Contain the virtual domain that should be used to represent a native VSAN in a specific VSAN (current VSAN). These 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). These 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

Note If you use persistent FC IDs for IVR, we recommend that you use them for all the devices in the IVR zoneset. We do not recommend using persistent FC IDs for some of the IVR devices while using automatic allocation for others.

Note IVR NAT must be enabled to use IVR persistent FC IDs.

Note 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.

Configuring Persistent FC IDs for IVR

To configure persistent FC IDs for IVR, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# ivr fcdomain database autonomous-fabric-num 21 vsan 22` `switch(config-fcdomain)#`	Enters IVR fcdomain database configuration submode for current AFID 21 and VSAN 22.
	`switch(config)# no ivr fcdomain database autonomous-fabric-num 21 vsan 22`	Deletes all the database entries, including all the corresponding persistent FC ID entries, for current AFID 21 and VSAN 22.
Step 3	`switch(config-fcdomain)# native-autonomous-fabric-num 20 native-vsan 11 domain 12` `switch(config-fcdomain-fcid)#`	Adds or replaces a database entry for native AFID 20, native VSAN 11, and domain 12, and enters IVR fcdomain FC ID configuration submode. Domains of all the corresponding persistent FC ID entries, if any, are also changed to 12.
	`switch(config-fcdomain)# no native-autonomous-fabric-num 20 native-vsan 11`	Deletes the virtual domain entry native AFID 20 and native VSAN 11, and all corresponding FC ID entries.
Step 4	`switch(config-fcdomain-fcid)# 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.
	`switch(config-fcdomain-fcid)# no pwwn 11:22:33:44:55:66:77:88`	Deletes the database entries for the pWWN.
Step 5	`switch(config-fcdomain-fcid)# device-alias SampleName fcid 0x123456`	Adds a database entry for mapping the device alias to the FC ID.
	`switch(config-fcdomain-fcid)# no device-alias SampleName`	Deletes the database entries for the device alias.

Verifying the Persistent FC ID Configuration

Verify the persistent FC ID configuration using the **show ivr fcdomain database** command. See Example 22-4 and Example 22-5

Example 22-4 Displays All IVR fcdomain Database Entries

```
switch# show ivr fcdomain database
-----------------------------------------------------
  AFID  Vsan  Native-AFID  Native-Vsan  Virtual-domain
-----------------------------------------------------
    1     2        10          11         0xc(12)
   21    22        20          11         0xc(12)

Number of Virtual-domain entries: 2

-----------------------------------------------------
  AFID  Vsan          Pwwn           Virtual-fcid
-----------------------------------------------------
   21    22  11:22:33:44:55:66:77:88  0x114466
   21    22  21:22:33:44:55:66:77:88  0x0c4466
   21    22  21:22:33:44:55:66:78:88  0x0c4466

Number of Virtual-fcid entries: 3
```

Example 22-5 Displays the IVR fcdomain Database Entries for a Specific AFID and VSAN

```
switch# show ivr fcdomain database autonomous-fabric-num 21 vsan 22
-------------------------------------------------
  AFID  Vsan  Native-AFID  Native-Vsan  Virtual-domain
-------------------------------------------------
   21    22        20          11         0xc(12)

Number of Virtual-domain entries: 1

-------------------------------------------------
  AFID  Vsan          Pwwn           Virtual-fcid
-------------------------------------------------
   21    22  11:22:33:44:55:66:77:88  0x114466
   21    22  21:22:33:44:55:66:77:88  0x0c4466
   21    22  21:22:33:44:55:66:78:88  0x0c4466

Number of Virtual-fcid entries: 3
```

Configuring IVR Logging Levels

To configure the severity level for logging messages from the IVR feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging level ivr 4`	Configures Telnet or SSH logging for the IVR feature at level 4 (warning). As a result, logging messages with a severity level of 4 or above are displayed.

Verifying Logging Level Configuration

Use the **show logging level** command to view the configured logging level for the IVR feature.

```
switch# show logging level
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
...
ivr                     5                       4
...
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

IVR Zones and IVR Zone Sets

As part of the IVR configuration, you need to configure one or more IVR zone 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.

Note The same IVR zone set must be activated on *all* of the IVR-enabled switches.

Caution Prior to Cisco SAN-OS Release 3.0(3) you can only configure a total of 10,000 zone members on all switches in a network. As of Cisco SAN-OS Release 3.0(3) you can only configure a total of 20,000 zone members on all switches in a network. A zone member is counted twice if it exists in two zones. See the "Database Merge Guidelines" section on page 22-37.

This section describes configuring IVR zones and IVR zone sets and includes the following topics:

About IVR Zones

Table 22-3 identifies the key differences between IVR zones and zones.

Table 22-3 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).

Automatic IVR Zone Creation

Figure 22-3 depicts an IVR zone consisting of four members. 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 22-3 Creating Zones Upon IVR Zone Activation

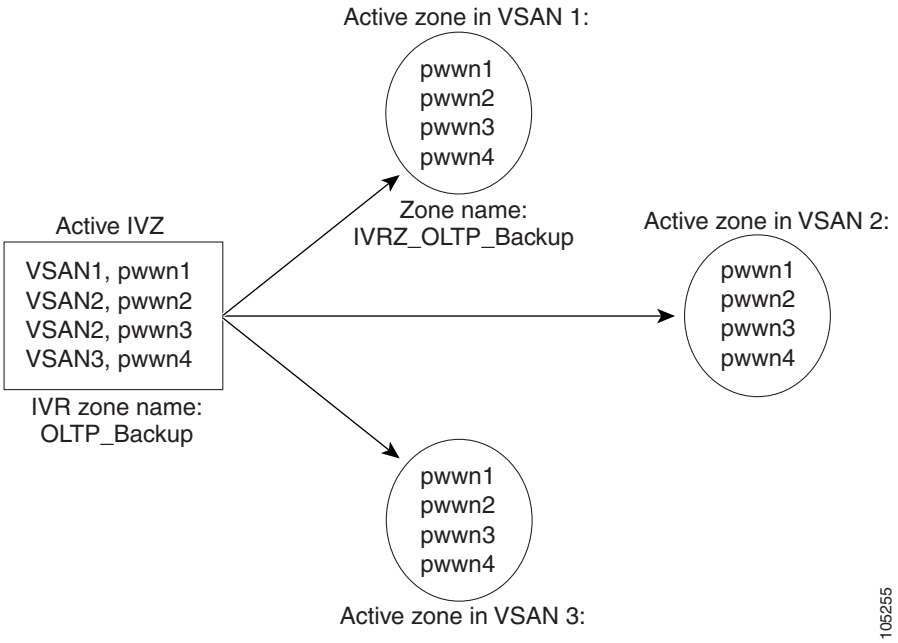

The zones are created automatically by the IVR process when an IVR zone set is activated. They are not stored in a full zone set database and are lost when the switch reboots or when a new zone set is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVR zone set configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated nondisruptively.

Note 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.

IVR zone and IVR zone set names are restricted to 64 alphanumeric characters.

Caution Prior to Cisco SAN-OS Release 3.0(3) you can only configure a total of 2000 IVR zones and 32 IVR zone sets on the switches in the network. As of Cisco SAN-OS Release 3.0(3) you can only configure a total of 8000 IVR zones and 32 IVR zone sets on the switches in the network. See the "Database Merge Guidelines" section on page 22-37.

Configuring IVR Zones and IVR Zone Sets

To create IVR zones and IVR zone sets, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr zone name sample_vsan2-3` `switch(config-ivr-zone)#`	Creates an IVR zone named sample_vsan2-3.

	Command	Purpose
Step 3	`switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:02:ca:4a vsan 3`	Adds the specified pWWN in VSAN 3 as an IVR zone member.
Step 4	`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.
Step 5	`switch(config-ivr-zone)# exit` `switch(config)#`	Reverts to configuration mode.
Step 6	`switch(config)# ivr zone name sample_vsan4-5` `switch(config-ivr-zone)#`	Creates an IVR zone named sample_vsan4-5.
Step 7	`switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:06:d9:1d vsan 4`	Adds the specified pWWN in VSAN 4 as an IVR zone member.
Step 8	`switch(config-ivr-zone)# member pwwn 21:01:00:e0:8b:2e:80:93 vsan 4`	Adds the specified pWWN in VSAN 4 as an IVR zone member.
Step 9	`switch(config-ivr-zone)# member pwwn 10:00:00:00:c9:2d:5a:dd vsan 5`	Adds the specified pWWN in VSAN 5 as an IVR zone member.
Step 10	`switch(config-ivr-zone)# exit` `switch(config)#`	Reverts to configuration mode.
Step 11	`switch(config)# ivr zoneset name Ivr_zoneset1` `switch(config-ivr-zoneset)#`	Creates an IVR zone set named Ivr_zoneset1.
Step 12	`switch(config-ivr-zoneset)# member sample_vsan2-3`	Adds the sample_vsan2-3 IVR zone as an IVR zone set member.
Step 13	`switch(config-ivr-zoneset)# member sample_vsan4-5`	Adds the sample_vsan4-5 IVR zone as an IVR zone set member.
Step 14	`switch(config-ivr-zoneset)# exit` `switch(config)`	Returns to configuration mode.
Step 15	`switch(config)# ivr zoneset activate name IVR_ZoneSet1`	Activates the newly created IVR zone set.
	`switch(config)# ivr zoneset activate name IVR_ZoneSet1 force`	Forcefully activates the specified IVR zone set.
	`switch(config)# no ivr zoneset activate name IVR_ZoneSet1`	Deactivates the specified IVR zone set.
Step 16	`switch(config)# end` `switch#`	Returns to EXEC mode.

About Activating Zone Sets and Using the force Option

Once the zone sets have been created and populated, you must activate the zone set. When you activate an IVR zone set, IVR automatically adds an IVR zone to the regular active zone set of each edge VSAN. If a VSAN does not have an active zone set, IVR can only activate an IVR zone set using the force option, which causes IVR to create an active zone set called “nozoneset” and adds the IVR zone to that active zone set.

Caution If you deactivate the regular active zone set in a VSAN, the IVR zone set is also deactivated. This occurs because the IVR zone in the regular active zone set, and all IVR traffic to and from the switch, is stopped. To reactivate the IVR zone set, you must reactivate the regular zone set.

Note If IVR and iSLB are enabled in the same fabric, at least one switch in the fabric must have both features enabled. Any zoning related configuration or activation operation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, traffic might be disrupted in the fabric.

You can also use the **force** option to activate IVR zone sets. Table 22-4 lists the various scenarios with and without the **force** option.

Table 22-4 IVR Scenarios with and without theforce Option.

Case	Default Zone Policy	Active Zone Set before IVR Zone Activation	`force` Option Used?	IVR Zone Set Activation Status	Active IVR Zone Created?	Possible Traffic Disruption
1	Deny	No active zone set	No	Failure	No	No
2			Yes	Success	Yes	No
3[1]	Deny	Active zone set present	No/Yes	Success	Yes	No
4	Permit	No active zone set *or* Active zone set present	No	Failure	No	No
5			Yes	Success	Yes	Yes

1. We recommend that you use the Case 3 scenario.

Caution Using the **force** option of IVR zone set activation may cause traffic disruption, even for devices that are not involved in IVR. For example, if your configuration does not have any active zone sets and the default zone policy is `permit`, then an IVR zone set activation will fail. However, IVR zone set activation will go through if the **force** option is used. Because zones are created in the edge VSANs corresponding to each IVR zone, traffic may be disrupted in edge VSANs where the default zone policy is `permit`.

Activating or Deactivating IVR Zone Sets

To activate or deactivate an existing IVR zone set, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr zoneset activate name IVR_ZoneSet1`	Activates the newly created IVR zone set.
	`switch(config)# ivr zoneset activate name IVR_ZoneSet1 force`	Forcefully activates the specified IVR zone set.
	`switch(config)# no ivr zoneset activate name IVR_ZoneSet1`	Deactivates the specified IVR zone set.

Note To replace the active IVR zone set with a new IVR zone set without disrupting traffic, activate the new IVR zone set without deactivating the current active IVR zone set.

Verifying IVR Zone and IVR Zone Set Configuration

Verify the IVR zone and IVR zone set configurations using the **show ivr zone** and **show ivr zoneset** commands. See Example 22-6 to Example 22-14.

Example 22-6 Displays the IVR Zone Configuration

```
switch# show ivr zone
zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2

zone name ivr_qa_z_all
    pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
    pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
    pwwn 10:00:00:00:c9:2d:5a:de vsan 2
    pwwn 21:00:00:20:37:5b:ce:af vsan 6
    pwwn 21:00:00:20:37:39:6b:dd vsan 6
    pwwn 22:00:00:20:37:39:6b:dd vsan 3
    pwwn 22:00:00:20:37:5b:ce:af vsan 3
    pwwn 50:06:04:82:bc:01:c3:84 vsan 5
```

Example 22-7 Displays Information for a Specified IVR Zone

```
switch# show ivr zone name sample_vsan2-3
zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 22-8 Displays the Specified Zone in the Active IVR Zone

```
switch# show ivr zone name sample_vsan2-3 active
zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 22-9 Displays the IVR Zone Set Configuration

```
switch# show ivr zoneset
zoneset name ivr_qa_zs_all
  zone name ivr_qa_z_all
    pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
    pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
    pwwn 10:00:00:00:c9:2d:5a:de vsan 2
    pwwn 21:00:00:20:37:5b:ce:af vsan 6
    pwwn 21:00:00:20:37:39:6b:dd vsan 6
    pwwn 22:00:00:20:37:39:6b:dd vsan 3
    pwwn 22:00:00:20:37:5b:ce:af vsan 3
    pwwn 50:06:04:82:bc:01:c3:84 vsan 5

zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 22-10 Displays the Active IVR Zone Set Configuration

```
switch# show ivr zoneset active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 22-11 Displays the Specified IVR Zone Set Configuration

```
switch# show ivr zoneset name IVR_ZoneSet1
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 22-12 Displays Brief Information for All IVR Zone Sets

```
switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
```

Example 22-13 Displays Brief Information for the Active IVR Zone Set

```
switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
```

Example 22-14 Displays Status Information for the IVR Zone Set

```
switch# show ivr zoneset status
Zoneset Status
_______________
    name               : IVR_ZoneSet1
    state              : activation success
    last activate time : Sat Mar 22 21:38:46 1980
    force option       : off

status per vsan:
_________________
     vsan     status
     ____     ______
       1      active
       2      active
```

Tip Repeat this configuration in all border switches participating in the IVR configuration.

Note Using the Cisco MDS Fabric Manager, you can distribute IVR zone configurations to all IVR-capable switches in the interconnected VSAN network. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

About 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. For more details on the advantages of LUN zoning, see the "About LUN Zoning" section on page 23-21.

Configuring LUNs in IVR Zoning

To configure LUNs in IVR zoning, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr zone name IvrLunZone` `switch(config-ivr-zone)#`	Configures an IVR zone called IvrLunZone.

	Command	Purpose
Step 3	`switch(config-ivr-zone)# member pwwn 10:00:00:23:45:67:89:ab lun 0x64 vsan 10`	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.
	`switch(config-ivr-zone)# member pwwn 10:00:00:23:45:67:89:ab lun 0x64 vsan 10 autonomous-fabric-id 20`	Configures an IVR zone member based on the specified pWWN, LUN value, and AFID.
	`switch(config-ivr-zone)# no member pwwn 20:81:00:0c:85:90:3e:80 lun 0x32 vsan 13 autonomous-fabric-id 10`	Removes an IVR zone member.

Note You can configure LUN zoning in an IVR zone set setup.

About QoS in IVR Zones

You can configure a QoS attribute for an IVR zone. The default QoS attribute setting is low.

Configuring the QoS Attribute

To configure the QoS attribute for an IVR zone, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ivr zone name IvrZone` `switch(config-ivr-zone)#`	Configures an IVR zone called IvrZone.
Step 3	`switch(config-ivr-zone)# attribute qos priority medium`	Configures the QoS for IVR zone traffic to medium.
	`switch(config-ivr-zone)# no attribute qos priority medium`	Reverts to the default QoS setting. The default is low.

Note If other QoS attributes are configured, the highest setting takes priority.

Verifying the QoS Attribute Configuration

Verify the QoS attribute configuration for an IVR zone using the **show ivr zone** command.

```
switch(config)# show ivr zone

zone name IvrZone
    attribute qos priority medium
```

Renaming IVR Zones and IVR Zone Sets

You can rename IVR zones and IVR zone sets.

To rename an IVR zone, use the **ivr zone rename** command in EXEC mode.

```
switch# ivr zone rename ivrzone1 ivrzone2
```

To rename an IVR zone set, use the **ivr zoneset rename** command in EXEC mode.

```
switch# ivr zoneset rename ivrzone1 ivrzone2
```

Clearing the IVR Zone Database

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

To clear the IVR zone database, use the **clear ivr zone database** command.

```
switch# clear ivr zone database
```

This command clears all configured IVR zone information.

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

Configuring IVR Using 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.

Note Read-only zoning cannot be configured in an IVR zone set setup.

System Image Downgrading Considerations

As of Cisco MDS SAN-OS Release 3.0(3), you can configure 8000 IVR zones and 20,000 IVR zone members. If you want to downgrade to a release prior to Cisco SAN-OS Release 3.0(3), the number of IVR zones cannot exceed 2000 and the number of IVR zone members cannot exceed 10,000.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the "CFS Merge Support" section on page 6-8 for detailed concepts.

- Be aware of the following conditions 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 (see Figure 22-4).

Figure 22-4 Fabric Merge Consequences

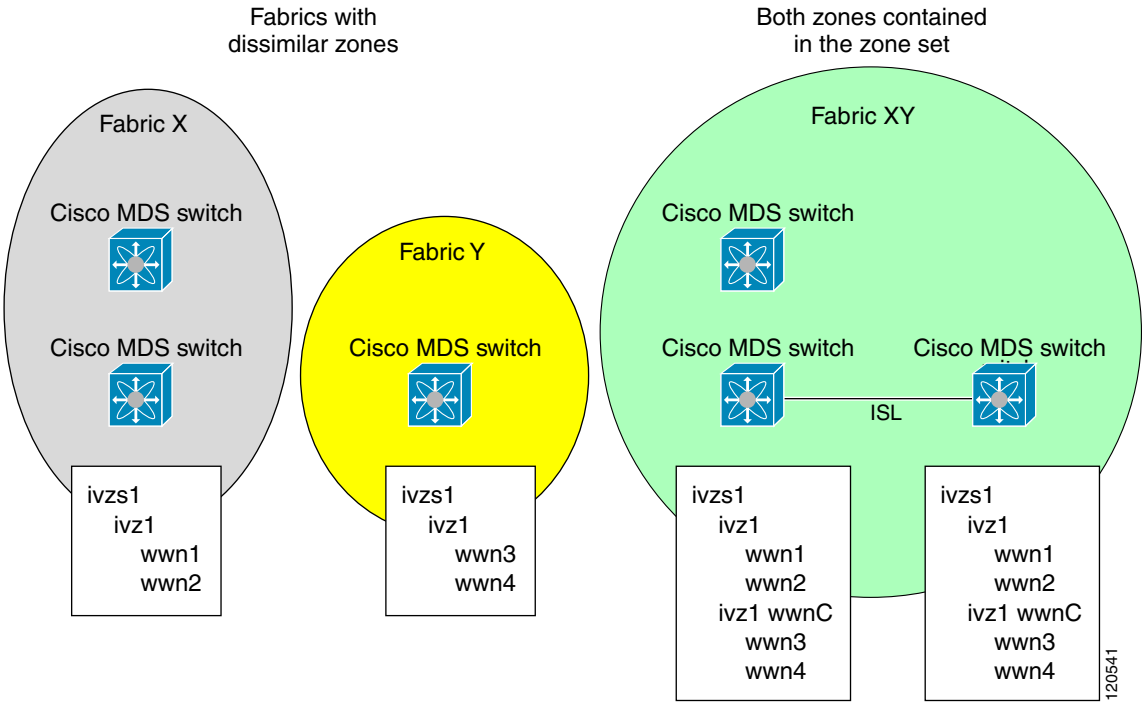

- You can configure different IVR configurations in different Cisco MDS switches.
- Be aware that the merge follows more liberal approach in order to avoid traffic disruption. After the merge, the configuration will be a union of the configurations that were present on the two switches involved in the merge.
  - The configurations are merged even if both fabrics have different configurations.
  - A union 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 union of the topology entries for both fabrics.
  - The merge will fail if the merged database contains more topology entries than the allowed maximum.
  - The total number of VSANs across the two fabrics cannot exceed 128.

Note VSANs with the same VSAN ID but different AFIDs are counted as two separate VSANs.

- The total number of IVR-enabled switches across the two fabrics cannot exceed 128.
- The total number of zone members across the two fabrics cannot exceed 10,000. As of Cisco SAN-OS Release 3.0(3), the total number of zone members across the two fabrics cannot exceed 20,000. A zone member is counted twice if it exists in two zones.

Note If only some of the switches in the fabrics are running Cisco SAN-OS Release 3.0(3) or later, and the number of zone members exceeds 10,000, you must either reduce the number of zone members in the fabric or upgrade all switches in both fabrics to Cisco SAN-OS Release 3.0(3) or later.

- The total number of zones across the two fabrics cannot exceed 2000. As of Cisco SAN-OS Release 3.0(3), the total number of zones across the two fabrics cannot exceed 8000.

Note If only some of the switches in the fabrics are running Cisco SAN-OS Release 3.0(3) or later, and if the number of zones exceeds 2000, you must either reduce the number of zones in the fabric or upgrade all switches in both fabrics to Cisco SAN-OS Release 3.0(3) or later.

- The total number or zone sets across the two fabrics cannot exceed 32.

Table 22-5 describes the results of a CFS merge of two IVR-enabled fabrics under different conditions.

Table 22-5 Results of Merging Two IVR-Enabled Fabrics

IVR Fabric 1	IVR Fabric 2	After Merge
NAT enabled	NAT disabled	Merge succeeds and NAT enabled
Auto mode on	Auto mode off	Merge succeeds and auto mode on
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.

Resolving Database Merge Failures

If a merge failure occurs, use the following commands to display the error conditions:

- **show ivr merge status**
- **show cfs merge status name ivr**
- **show logging last** *lines* (and look for MERGE failures)

Depending on the failure indicated in the **show** command outputs, you can perform the following:

- If the failure is due to exceeding the maximum configuration limits in a fabric where the switches are running more than one Cisco SAN-OS release, then either upgrade the switches running the earlier release or reduce the number of IVR zones and IIVR zone members on the switches running the more recent release to the earlier release limit (see the"IVR Limits Summary" section on page 22-4).
- If the failure is due to exceeding maximum limits in a fabric where all switches are running the same Cisco SAN-OS release, identify the switch that has the correct configuration and perform a CFS commit to distribute the IVR configuration (see the "Configuring Default AFIDs" section on page 22-16 and the"IVR Limits Summary" section on page 22-4).
- 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 (see the "Configuring Individual AFIDs" section on page 22-17).

After a successful CFS commit, the merge will be successful.

Example Configurations

This section provides IVR configurations examples and includes the following topics:

Manual Topology Configuration

This section provides the configuration steps to manually configure the example illustrated in Figure 22-1.

Step 1 Enable IVR.

```
mds# config t
Enter configuration commands, one per line.  End with CNTL/Z.
mds(config)# ivr enable
mds(config)# exit
mds#
```

Step 2 Verify that IVR is enabled.

```
mds# 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.

```
mds# config t
Enter configuration commands, one per line.  End with CNTL/Z.
mds(config)# ivr distribution
```

Step 4 Manually configure the IVR VSAN-topology. In Figure 22-1, two of the four IVR-enabled switches (MDS1 and MDS2) are members of VSANs 1 and 4. The other two switches (MDS3 and MDS4) are members of VSANs 2, 3, and 4.

```
mds(config)# ivr vsan-topology database
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:05:40:01:1b:c2
vsan-ranges 1,4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:02:00:44:22:00:4a:08
vsan-ranges 1,4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:02:8a:04
vsan-ranges 2-4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:40:aa:16
vsan-ranges 2-4
mds(config-ivr-topology-db)# exit
mds(config)#
```

Step 5 Verify the configured VSAN topology.

Note The configured topology has not yet been activated—as indicated by the `no` status displayed in the `Active` column.

```
mds(config)# do show ivr vsan-topology

AFID  SWITCH WWN                 Active   Cfg. VSANS
-------------------------------------------------------------
   1  20:00:00:05:40:01:1b:c2 *    no     yes  1,4
   1  20:00:00:44:22:00:4a:08      no     yes  1,4
   1  20:00:00:44:22:02:8a:04      no     yes  2-4
   1  20:00:00:44:22:40:aa:16      no     yes  2-4
```

```
Total:   4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE
```

Step 6 Activate the configured VSAN topology.

```
mds(config)# ivr vsan-topology activate
```

Step 7 Verify the activation.

```
mds(config)# do show ivr vsan-topology

AFID  SWITCH WWN                 Active   Cfg. VSANS
------------------------------------------------------------
   1  20:00:00:05:40:01:1b:c2 *   yes     yes  1,4
   1  20:00:00:44:22:00:4a:08     yes     yes  1,4
   1  20:00:00:44:22:02:8a:04     yes     yes  2-4
   1  20:00:00:44:22:40:aa:16     yes     yes  2-4

Total:   4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Tue May 20 23:14:59 1980
```

Step 8 Configure IVR zone set and zones. 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).

Tip Instead of creating two IVR zones, you can also create one IVR zone with the tape and both servers.

```
mds(config)# ivr zoneset name tape_server1_server2

mds(config-ivr-zoneset)# zone name tape_server1
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
mds(config-ivr-zoneset-zone)# member pwwn 10:02:66:45:00:20:89:04 vsan 2
mds(config-ivr-zoneset-zone)# exit

mds(config-ivr-zoneset)# zone name tape_server2
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
mds(config-ivr-zoneset-zone)# member pwwn 10:00:ad:51:78:33:f9:86 vsan 3
mds(config-ivr-zoneset-zone)# exit
```

Step 9 View the IVR zone configuration to confirm that the IVR zone set and IVR zones are properly configured.

```
mds(config)# do 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.

```
mds(config)# do 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.

```
mds(config)# ivr zoneset activate name tape_server1_server2
zoneset activation initiated. check inter-VSAN zoneset status
mds(config)# exit
mds#
```

Step 12 Verify the IVR zone set activation.

```
mds# 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.

```
mds# 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 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

mds# 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
```

Auto-Topology Configuration

This section provides example configuration steps for configuring IVR auto-topology.

Step 1 Enable IVR on every border switch in the fabric.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# ivr enable
switch(config)# exit
switch#
```

Step 2 Verify that IVR is enabled on every IVR-enabled switch.

```
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.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# ivr distribution
```

Step 4 Enable IVR auto-topology mode.

```
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.

```
switch(config)# ivr commit
switch(config)# exit
switch#
```

Step 6 Verify the status of the commit request.

```
switch# show ivr session status
Last Action                : Commit
Last Action Result         : Success
Last Action Failure Reason : None
```

Step 7 Verify the active IVR topology.

```
switch# show ivr vsan-topology active

AFID  SWITCH WWN                 Active   Cfg. VSANS
-----------------------------------------------------------------
   1  20:00:00:0d:ec:08:6e:40 *  yes      no   1,336-338
   1  20:00:00:0d:ec:0c:99:40    yes      no   336,339
```

Default Settings

Table 22-6 lists the default settings for IVR parameters.

Table 22-6 Default IVR Parameters

Parameters	Default
IVR feature	Disabled.
IVR VSANs	Not added to virtual domains.
IVR NAT	Disabled.
QoS for IVR zones	Low.
Configuration distribution	Disabled.

CHAPTER 23

Configuring and Managing Zones

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

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

This chapter includes the following sections:

Note Table 19-1 on page 19-4 lists the differences between zones and VSANs.

About Zoning

Zoning has the following features:

- A zone consists of multiple zone members.
  - Members in a zone can access each other; members in different zones cannot access each other.
  - If zoning is not activated, all devices are members of the default zone.
  - If zoning is activated, any device that is not in an active zone (a zone that is part of an active 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 criteria is based mainly on WWNs or FC IDs.
  - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
  - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
  - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
  - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
  - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
  - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
  - IPv4 address—Specifies the IPv4 address (and optionally the subnet mask) of an attached device.

  - IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.

- 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

Figure 23-1 illustrates 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. Note that H3 resides in both zones.

Figure 23-1 Fabric with Two Zones

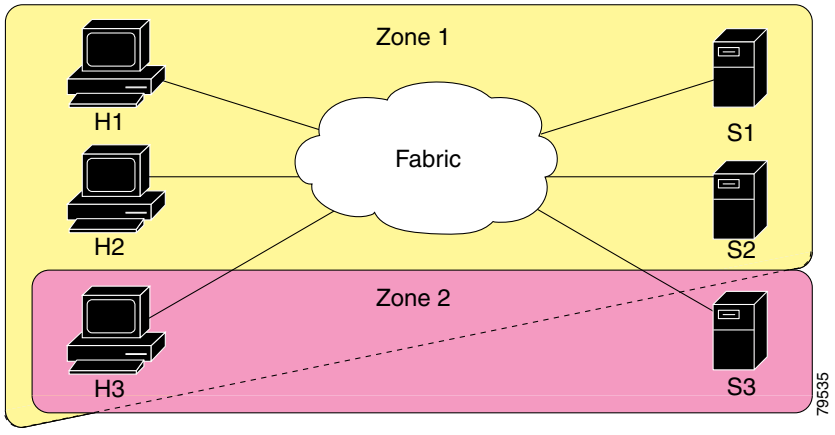

Of course, there are other ways to partition this fabric into zones. Figure 23-2 illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 23-2 Fabric with Three Zones

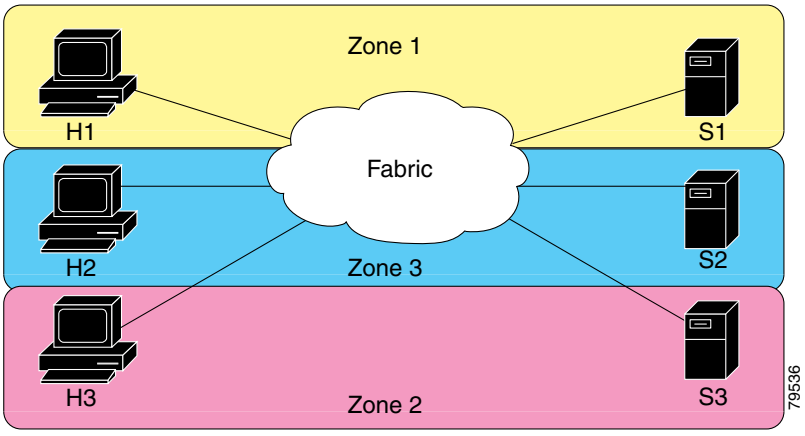

Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

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

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

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

Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

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

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

Figure 23-3 shows a zone being added to an activated zone set.

Figure 23-3 Active and Full Zone Sets

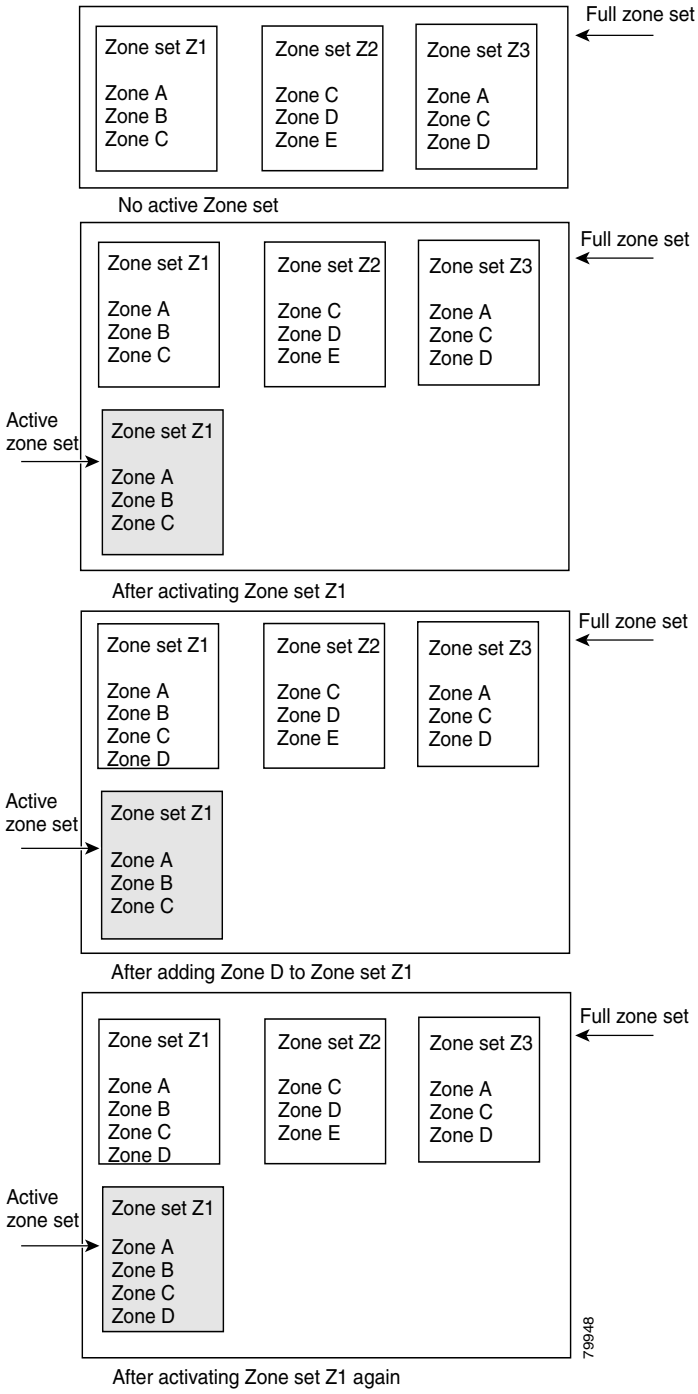

Zone Configuration

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

Configuring a Zone

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zone name Zone1 vsan 3` `switch(config-zone)#`	Configures a zone called Zone1 for the VSAN called vsan3. **Note** All alphanumeric characters or one of the following symbols ($, -, ^, _) are supported.

Step 3	`switch(config-zone)# member` *type value*
	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`
	IPv4 address example:
	`switch(config-zone)# member ip-address 10.15.0.0 255.255.0.0`
	IPv6 address example:
	`switch(config-zone)# member ipv6-address 2001::db8:800:200c:417a/64`
	Local sWWN interface example:
	`switch(config-zone)# member interface fc 2/1`
	Remote sWWN interface example:
	`switch(config-zone)# member interface fc2/1 swwn 20:00:00:05:30:00:4a:de`
	Domain ID interface example:
	`switch(config-zone)# member interface fc2/1 domain-id 25`
		Configures a member for the specified zone (Zone1) based on the type (pWWN, fabric pWWN, FC ID, fcalias, domain ID, IPv4 address, IPv6 address, or interface) and value specified.
		Caution You must only configure pWWN-type zoning on all MDS switches running Cisco SAN-OS if there is a Cisco MDS 9020 switch running FabricWare in the same fabric.

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

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

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

Zone Sets

In Figure 23-4, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 23-4 Hierarchy of Zone Sets, Zones, and Zone Members

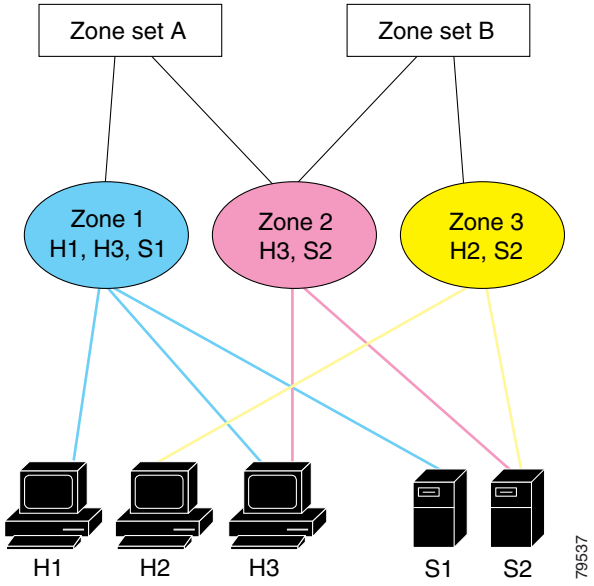

Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

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

Activating a Zone Set

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

To activate or deactivate an existing zone set, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zoneset activate name Zoneset1 vsan 3`	Activates the specified zone set.
	`switch(config)# no zoneset activate name Zoneset1 vsan 3`	Deactivates the specified zone set

About the Default Zone

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

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

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

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

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

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

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

Configuring the Default Zone Access Permission

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# zone default-zone permit vsan 1`	Permits traffic flow to default zone members.
	`switch(config)# no zone default-zone permit vsan 1`	Denies (default) traffic flow to default zone members.

About FC Alias Creation

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

- pWWN—The WWN of the N or NL port is in hex format (for example, 10:00:00:23:45:67:89:ab).
- fWWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- IPv4 address—The IPv4 address of an attached device is in 32 bits in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv6 address—The IPv6 address of an attached device is in 128 bits in colon- (:) separated) hexadecimal format.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

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

Creating FC Aliases

To create an alias, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcalias name AliasSample vsan 3` `switch(config-fcalias)#`	Configures an alias name (AliasSample).

	Command	Purpose
Step 3	switch(config-fcalias)# **member** *type value* pWWN example: switch(config-fcalias)# **member pwwn 10:00:00:23:45:67:89:ab** fWWN example: switch(config-fcalias)# **member fwwn 10:01:10:01:10:ab:cd:ef** FC ID example: switch(config-fcalias)# **member fcid 0x222222** Domain ID example: switch(config-fcalias)# **member domain-id 2 portnumber 23** IPv4 address example: switch(config-fcalias)# **member ip-address 10.15.0.0 255.255.0.0** IPv6 address example: switch(config-fcalias)# **member ipv6-address 2001::db8:800:200c:417a/64** Local sWWN interface example: switch(config-fcalias)# **member interface fc 2/1** Remote sWWN interface example: switch(config-fcalias)# **member interface fc2/1 swwn 20:00:00:05:30:00:4a:de** Domain ID interface example: switch(config-fcalias)# **member interface fc2/1 domain-id 25**	Configures a member for the specified fcalias (AliasSample) based on the type (pWWN, fabric pWWN, FC ID, domain ID, IPv4 address, IPv6 address, or interface) and value specified.
Step 4	**Note** Multiple members can be specified on multiple lines.	

Creating Zone Sets and Adding Member Zones

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

	Command	Purpose
Step 1	switch# **config t**	Enters configuration mode.
Step 2	switch(config)# **zoneset name Zoneset1 vsan 3** switch(config-zoneset)#	Configures a zone set called Zoneset1. **Tip** To activate a zone set, you must first create the zone and a zone set.
Step 3	switch(config-zoneset)# **member Zone1**	Adds Zone1 as a member of the specified zone set (Zoneset1). **Tip** If the specified zone name was not previously configured, this command will return the `Zone not present` error message.

	Command	Purpose
Step 4	`switch(config-zoneset)# zone name InlineZone1` `switch(config-zoneset-zone)#`	Adds a zone (InlineZone1) to the specified zone set (Zoneset1). **Tip** Execute this step only if you need to create a zone from a zone set prompt.
Step 5	`switch(config-zoneset-zone)# member fcid 0x111112` `switch(config-zoneset-zone)#`	Adds a new member (FC ID 0x111112) to the new zone (InlineZone1). **Tip** Execute this step only if you need to add a member to a zone from a zone set prompt.

Tip You do not have to issue the **copy running-config startup-config** command to store the active zone set. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. It is not available across switch resets.

Caution If you deactivate the active zone set in a VSAN that is also configured for IVR, the active IVR zone set (IVZS) is also deactivated and all IVR traffic to and from the switch is stopped. This deactivation can disrupt traffic in more than one VSAN. Before deactivating the active zone set, check the active zone analysis for the VSAN (see the "Zone and Zone Set Analysis" section on page 23-41). To reactivate the IVZS, you must reactivate the regular zone set (see the "Configuring IVR Zones and IVR Zone Sets" section on page 22-29).

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

Note The pWWN of the virtual target does not appear in the zoning end devices database in Fabric Manager. If you want to zone the virtual device with a pWWN, you must enter it in the Add Member to Zone dialog box when creating a zone. However, if the device alias is in enhanced mode, the virtual device names appear in the device alias database in the Fabric Manager zoning window. In this case, users can choose to select either the device alias name or enter the pWWN in the Add Member to Zone dialog box.

Note Set the device alias mode to **enhanced** when using SDV (because the pWWN of a virtual device could change).

For example, SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenable SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. Hence, you would have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or

when the pWWN changes.

Be sure you understand how device alias modes work before enabling them. Refer to Chapter 24, "Distributing Device Alias Services" for details and requirements about device alias modes.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the 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 Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.

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

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution at the EXEC mode level or full zone set distribution at the configuration mode level. Table 23-1 lists the differences.

Table 23-1 zoneset distribution Command Differences

zoneset distribute vsan Command (EXEC Mode)	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.

This section describes zone set distribution and includes the following topics:

- Enabling Full Zone Set Distribution, page 23-14
- Enabling a One-Time Distribution, page 23-14
- About Recovering from Link Isolation, page 23-15
- Importing and Exporting Zone Sets, page 23-15

Enabling Full Zone Set Distribution

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

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zoneset distribute full vsan 33`	Enables sending a full zone set along with an active zone set.

Enabling a One-Time Distribution

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

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

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

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

Note The **zoneset distribute vsan** *vsan-id* command is supported in **interop 2** and **interop 3** modes—not in **interop 1** mode.

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

```
switch# show zone status vsan 2
VSAN: 3 default-zone: permit distribute: active only Interop: 100
    mode:basic merge-control:allow session:none
    hard-zoning:enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:0  Zones:0 Aliases: 0
Active Zoning Database :
    Name: nozoneset  Zonesets:1  Zones:2
Status: Zoneset distribution completed at 04:01:06 Aug 28 2004
```

About Recovering from Link Isolation

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

- Import the neighboring switch's active zone set database and replace the current active zone set (see Figure 23-5).
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 23-5 Importing and Exporting the Database

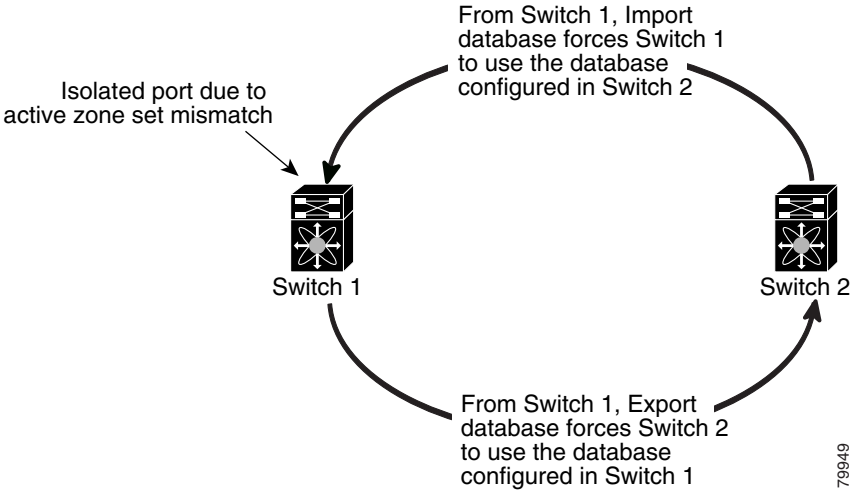

Importing and Exporting Zone Sets

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

	Command	Purpose
Step 1	`switch# zoneset import interface fc1/3 vsan 2`	Imports the zone set from the adjacent switch connected through the fc 1/3 interface for VSAN 2.
	`switch# zoneset import interface fc1/3 vsan 2-5`	Imports the zone set from the adjacent switch connected through the fc 1/3 interface for VSANs ranging from 2 through 5.
Step 2	`switch# zoneset export vsan 5`	Exports the zone set to the adjacent switch connected through VSAN 5.
	`switch# zoneset export vsan 5-8`	Exports the zone set to the adjacent switch connected through the range of VSANs 5 through 8.

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

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.

This section includes the following topics:

Copying Zone Sets

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

To make a copy of a zone set, follow this step:

	Command	Purpose
Step 1	`switch# zone copy active-zoneset full-zoneset vsan 2` `Please enter yes to proceed.(y/n) [n]? y`	Makes a copy of the active zone set in VSAN 2 to the full zone set.
	`switch# zone copy vsan 3 active-zoneset scp://guest@myserver/tmp/active_zoneset.txt`	Copies the active zone in VSAN 3 to a remote location using SCP.

Caution If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zone set, then a zone set copy operation copies all the IVR zones to the full zone database. To prevent copying to the IVR zones, you must explicitly remove them from the full zone set database before performing the copy operation. Refer to the Chapter 22, "Configuring Inter-VSAN Routing" for more information on the IVR feature.

Renaming Zones, Zone Sets, and Aliases

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zoneset rename oldname newname vsan 2`	Renames a zone set in the specified VSAN.
	`switch(config)# zone rename oldname newname vsan 2`	Renames a zone in the specified VSAN.
	`switch(config)# fcalias rename oldname newname vsan 2`	Renames a fcalias in the specified VSAN.
	`switch(config)# zone-attribute-group rename oldname newname vsan 2`	Renames a zone attribute group in the specified VSAN.
Step 3	`switch(config)# zoneset activate name newname vsan 2`	Activates the zone set and updates the new zone name in the active zone set.

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

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zoneset clone oldname newname vsan 2`	Clones a zone set in the specified VSAN.
	`switch(config)# zone clone oldname newname vsan 2`	Clones a zone in the specified VSAN.
	`switch(config)# fcalias clone oldname newname vsan 2`	Clones a fcalias in the specified VSAN.
	`switch(config)# zone-attribute-group clone oldname newname vsan 2`	Clones a zone attribute group in the specified VSAN.
Step 3	`switch(config)# zoneset activate name newname vsan 2`	Activates the zone set and updates the new zone name in the active zone set.

Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN. To clear the zone server database, use the following command:

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

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

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

Advanced Zone Attributes

This section describes advanced zone attributes and includes the following topics:

About Zone-Based Traffic Priority

The zoning feature provides an additional segregation mechanism to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the Quality of Service (QoS) priority as a zone attribute. You can assign the QoS traffic priority attribute to be high, medium, or low. By default, zones with no specified priority are implicitly assigned a low priority. See the "VSAN Versus Zone-Based QoS" section on page 56-7 for more information.

To use this feature, you need to obtain the ENTERPRISE_PKG license (see Chapter 3, "Obtaining and Installing Licenses") and you must enable QoS in the switch (see the "About Data Traffic" section on page 56-6).

This feature allows SAN administrators to configure QoS in terms of a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis rather than between zone members.

Caution If zone-based QoS is implemented in a switch, you cannot configure the interop mode in that VSAN.

Configuring Zone-Based Traffic Priority

To configure the zone priority, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zone name QosZone vsan 2` `switch(config-zone)#`	Configures an alias name (QosZone) and enters zone configuration submode.

	Command	Purpose
Step 3	`switch(config-zone)# attribute qos priority high`	Configures this zone to assign high priority QoS traffic to each frame matching this zone.
	`switch(config-zone)# attribute qos priority medium`	Configures this zone to assign medium priority QoS traffic to each frame matching this zone.
	`switch(config-zone)# attribute qos priority low`	Configures this zone to assign low priority QoS traffic to each frame matching this zone.
	`switch(config-zone)# no attribute qos priority high`	Reverts to using the default low priority for this zone.
Step 4	`switch(config-zone)# exit` `switch(config)#`	Returns to configuration mode.
Step 5	`switch(config)# zoneset name QosZoneset vsan 2` `switch(config-zoneset)#`	Configures a zone set called QosZoneset for the specified VSAN (vsan 2) and enters zone set configuration submode. **Tip** To activate a zone set, you must first create the zone and a zone set.
Step 6	`switch(config-zoneset)# member QosZone`	Adds QosZone as a member of the specified zone set (QosZoneset). **Tip** If the specified zone name was not previously configured, this command will return the `Zone not present` error message.
Step 7	`switch(config-zoneset)# exit` `switch(config)#`	Returns to configuration mode.
Step 8	`switch(config)# zoneset activate name QosZoneset vsan 2`	Activates the specified zone set.

Configuring Default Zone QoS Priority Attributes

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

Note If a member is part of two zones with two different QoS priority attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone default-zone vsan 1` `switch(config-default-zone)#`	Enters the default zone configuration submode.

	Command	Purpose
Step 3	`switch(config-default-zone)# attribute qos priority high`	Sets the QoS priority attribute for frames matching these zones.
	`switch(config-default-zone)# no attribute qos priority high`	Removes the QoS priority attribute for the default zone and reverts to default low priority.

About Broadcast Zoning

Note Broadcast zoning is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

You can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled and broadcast frames are sent to all Nx ports in the VSAN. When enabled, broadcast frames are only sent to Nx ports in the same zone, or zones, as the sender. Enable broadcast zoning when a host or storage device uses this feature.

Tip If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

Caution If broadcast zoning is enabled on a switch, you cannot configure the interop mode in that VSAN.

Configuring Broadcast Zoning

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone broadcast enable vsan 2`	Broadcasts frames for the specified VSAN.
	`switch(config)# no zone broadcast enable vsan 3`	Disables (default) broadcasting for the specified VSAN.
Step 3	`switch(config)# zone name BcastZone vsan 2` `switch(config-zone)#`	Creates a broadcast zone in the specified VSAN and enters zone configuration submode.
Step 4	`switch(config-zone)# member pwwn 21:00:00:20:37:f0:2e:4d`	Adds the specified member to this zone.
Step 5	`switch(config-zone)# attribute broadcast`	Specifies this zone to be broadcast to other devices.

Step 6	`switch(config-zone)# end`
	`switch# show zone vsan 2`
	`zone name bcast-zone vsan 2`
	`attribute broadcast`
	`pwwn 21:00:00:e0:8b:0b:66:56`
	`pwwn 21:00:00:20:37:f0:2e:4d`
		Displays the broadcast configuration

To configure the **broadcast** attribute for a default zone, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone default-zone vsan 1` `switch(config-default-zone)#`	Enters the default zone configuration submode.
Step 3	`switch(config-default-zone)# attribute broadcast`	Sets broadcast attributes for the default zone.
	`switch(config-default-zone)# no attribute broadcast`	Reverts the default zone attributes to read-write (default).

About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

Caution LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

Note When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.

- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

Note Unzoned LUNs automatically become members of the default zone.

Figure 23-6 shows a LUN-based zone example.

Figure 23-6 LUN Zoning Access

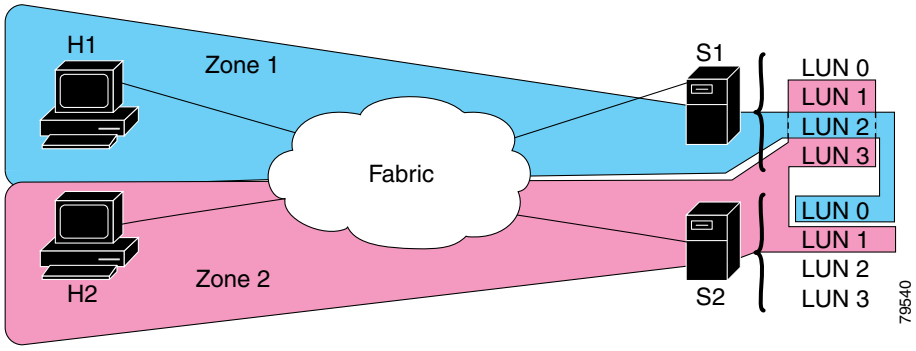

Configuring a LUN-Based Zone

To configure a LUN-based zone, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone name LunSample vsan 2` `switch(config-zone)#`	Configures a zone called LunSample for the specified VSAN (vsan 2) and enters zone configuration submode.
Step 3	`switch(config-zone)# member pwwn 10:00:00:23:45:67:89:ab lun 0x64`	Configures a 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. LUN 0x64 in hex format corresponds to 100 in decimal format.
	`switch(config-zone)# member fcid 0x12465 lun 0x64`	Configures a zone member based on the FC ID and LUN value.

Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each host bus adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the "Configuring a LUN-Based Zone" section on page 23-22.

Note Refer to the relevant user manuals to obtain the LUN number for each HBA.

Caution If you make any errors when assigning LUNs, you might lose data.

About Read-Only Zones

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones.

Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, the read-only zone takes priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.

  The read-only zone feature behaves as designed if either the FAT16 or FAT32 file system is used with the previously mentioned Windows operating systems.

Configuring Read-Only Zones

To configure read-only zones, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone name Sample2 vsan 2` `switch(config-zone)#`	Configures a zone called Sample2 for the specified VSAN (vsan 2) and enters zone configuration submode.
Step 3	`switch(config-zone)# attribute read-only`	Sets read-only attributes for the Sample2 zone. **Note** The default is read-write for all zones.
	`switch(config-zone)# no attribute read-only`	Reverts the Sample2 zone attributes to read-write.

To configure the **read-only** option for a default zone, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone default-zone vsan 1` `switch(config-default-zone)#`	Enters the default zone configuration submode.
Step 3	`switch(config-default-zone)# attribute read-only`	Sets read-only attributes for the default zone.
	`switch(config-default-zone)# no attribute read-only`	Reverts the default zone attributes to read-write (default).

Displaying Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed. See Examples 23-1 to 23-16.

Example 23-1 Displays Zone Information for All VSANs

```
switch# show zone
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 2
  fwwn 20:41:00:05:30:00:2a:1e
  fwwn 20:42:00:05:30:00:2a:1e
  fwwn 20:43:00:05:30:00:2a:1e

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

zone name Techdocs vsan 3
  ip-address 10.15.0.0 255.255.255.0

zone name Zone21 vsan 5
  pwwn 21:00:00:20:37:a6:be:35
  pwwn 21:00:00:20:37:a6:be:39
  fcid 0xe000ef
  fcid 0xe000e0
  symbolic-nodename iqn.test
  fwwn 20:1f:00:05:30:00:e5:c6
  fwwn 12:12:11:12:11:12:12:10
  interface fc1/5 swwn 20:00:00:05:30:00:2a:1e
  ip-address 12.2.4.5 255.255.255.0
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35

zone name Zone2 vsan 11
  interface fc1/5 pwwn 20:4f:00:05:30:00:2a:1e
```

```
zone name Zone22 vsan 6
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35

zone name Zone23 vsan 61
  pwwn 21:00:00:04:cf:fb:3e:7b lun 0000
```

Example 23-2 Displays Zone Information for a Specific VSAN

```
switch# show zone vsan 1
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 1
    fwwn 20:4f:00:05:30:00:2a:1e
    fwwn 20:50:00:05:30:00:2a:1e
    fwwn 20:51:00:05:30:00:2a:1e
    fwwn 20:52:00:05:30:00:2a:1e
    fwwn 20:53:00:05:30:00:2a:1e

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

Use the **show zoneset** command to view the configured zone sets.

Example 23-3 Displays Configured Zone Set Information

```
switch# show zoneset vsan 1
zoneset name ZoneSet2 vsan 1
  zone name Zone2 vsan 1
    fwwn 20:4e:00:05:30:00:2a:1e
    fwwn 20:4f:00:05:30:00:2a:1e
    fwwn 20:50:00:05:30:00:2a:1e
    fwwn 20:51:00:05:30:00:2a:1e
    fwwn 20:52:00:05:30:00:2a:1e

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

zoneset name ZoneSet1 vsan 1
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

Example 23-4 Displays Configured Zone Set Information for a Range of VSANs

```
switch# show zoneset vsan 2-3
zoneset name ZoneSet2 vsan 2
  zone name Zone2 vsan 2
    fwwn 20:52:00:05:30:00:2a:1e
    fwwn 20:53:00:05:30:00:2a:1e
    fwwn 20:54:00:05:30:00:2a:1e
```

```
    fwwn 20:55:00:05:30:00:2a:1e
    fwwn 20:56:00:05:30:00:2a:1e

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

zoneset name ZoneSet3 vsan 3
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

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

Example 23-5 Displays Members of a Zone

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

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

Example 23-6 Displays fcalias Configuration

```
switch# show fcalias vsan 1
fcalias name Alias2 vsan 1

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

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

Example 23-7 Displays Membership Status

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

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

Example 23-8 Displays Zone Statistics

```
switch# show zone statistics
Statistics For VSAN: 1
********************************
Number of Merge Requests Sent: 24
Number of Merge Requests Recvd: 25
Number of Merge Accepts Sent: 25
Number of Merge Accepts Recvd: 25
```

```
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
Statistics For VSAN: 2
********************************
Number of Merge Requests Sent: 4
Number of Merge Requests Recvd: 4
Number of Merge Accepts Sent: 4
Number of Merge Accepts Recvd: 4
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
```

Example 23-9 Displays LUN Zone Statistics

```
switch# show zone statistics lun-zoning
LUN zoning statistics for VSAN: 1
************************************************************
S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00:00
------------------------------------------------------------
Number of Inquiry commands received:            10
Number of Inquiry data No LU sent:               5
Number of Report LUNs commands received:        10
Number of Request Sense commands received:       1
Number of Other commands received:               0
Number of Illegal Request Check Condition sent:  0

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

Example 23-10 Displays LUN Zone Statistics

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

Example 23-11 Displays Active Zone Sets

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

Example 23-12 Displays Brief Descriptions of Zone Sets

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

Example 23-13 Displays Active Zones

```
switch# show zone active
zone name Zone2 vsan 1
* fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

zone name IVRZ_IvrZone1 vsan 1
  pwwn 10:00:00:00:77:99:7a:1b
* fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]

zone name IVRZ_IvrZone4 vsan 1
* fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
* fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

zone name Zone1 vsan 1667
  fcid 0x123456

zone name $default_zone$ vsan 1667
```

Example 23-14 Displays Active Zone Sets

```
switch# show zoneset active
zoneset name ZoneSet4 vsan 1
  zone name Zone2 vsan 1
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

  zone name IVRZ_IvrZone1 vsan 1
    pwwn 10:00:00:00:77:99:7a:1b
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]

zoneset name QosZoneset vsan 2
  zone name QosZone vsan 2
  attribute qos priority high
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]
```

```
Active zoneset vsan 1667
  zone name Zone1 vsan 1667
    fcid 0x123456

  zone name $default_zone$ vsan 1667
```

Example 23-15 Displays Zone Status

```
switch# show zone status
VSAN: 1 default-zone: deny distribute: full Interop: Off
    mode:basic merge-control:allow session:none
    hard-zoning:enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:1  Zones:11  Aliases:0
Active Zoning Database :
    Name: zoneset-1  Zonesets:1  Zones:11  Aliases:0
Status: Activation completed at Thu Feb 13 10:22:34 2003

VSAN: 2 default-zone: deny distribute: full Interop: Off
    mode:basic merge-control:allow session:none
    hard-zoning:enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:1  Zones:10  Aliases:0
Active Zoning Database :
    Name: zoneset-2  Zonesets:1  Zones:10  Aliases:0
Status: Activation completed at Thu Feb 13 10:23:12 2003

VSAN: 3 default-zone: deny distribute: full Interop: Off
    mode:basic merge-control:allow session:none
    hard-zoning:enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:1  Zones:10  Aliases:0
Active Zoning Database :
    Name: zoneset-3  Zonesets:1  Zones:10  Aliases:0
Status: Activation completed at Thu Feb 13 10:23:50 2003
```

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

Example 23-16 Displays Zone Statistics

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

Use the **show running** and **show zone active** commands to display the configured interface-based zones (see Example 23-17 and Example 23-18).

Example 23-17 Displays the Interface-Based Zones

```
switch# show running
zone name if-zone vsan 1
      member interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2
      member fwwn 20:4f:00:0c:88:00:4a:e2
      member interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
      member pwwn 22:00:00:20:37:39:6b:dd
```

Example 23-18 Displays the fWWNs and Interfaces in an Active Zone

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

A similar output is also available on the remote switch (see Example 23-19).

Example 23-19 Displays the Local Interface Active Zone Details for a Remote Switch

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

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.

This section includes the following topics:

About Enhanced Zoning

Table 23-2 lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Family.

Table 23-2 Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.
To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.
The MDS-specific zone member types (IPv4 address, IPv6 address, symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the MDS-specific types can be misunderstood by the non-Cisco switches.	Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type.	Unique vendor type.
The fWWN-based zone membership is only supported in Cisco interop mode.	Supports fWWN-based membership in the standard interop mode (interop mode 1).	The fWWN-based member type is standardized.

Changing from Basic Zoning to Enhanced Zoning

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

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

If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.

Step 2 Set the operation mode to enhanced zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies and then release the lock. All switches in the fabric then move to the enhanced zoning mode.

Tip After moving from basic zoning to enhanced zoning we recommend that you save the running configuration.

Changing from Enhanced Zoning to Basic Zoning

The standards do not allow you to move back to basic zoning. However, Cisco MDS switches allow this move to enable you to downgrade and upgrade to other Cisco SAN-OS releases.

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

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

If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the Cisco SAN-OS software automatically removes them.

Step 2 Set the operation mode to basic zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.

Note If a switch running Cisco SAN-OS Release 2.0(1b), or later, with enhanced zoning enabled is downgraded to Cisco SAN-OS Release 1.3(4), or earlier, the switch comes up in basic zoning mode and thus cannot join the fabric because all the other switches in the fabric are still in enhanced zoning mode.

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family.

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone mode enhanced vsan 3000` `Set zoning mode command initiated. Check zone status`	Enables enhanced zoning in the specified VSAN.
	`switch(config)# no zone mode enhanced vsan 150` `Set zoning mode command initiated. Check zone status`	Disables enhanced zoning in the specified VSAN.

Modifying the Zone Database

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

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

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone commit vsan 2` `No pending info found`	Applies the changes to the enhanced zone database and closes the session.
	`switch(config)# zone commit vsan 3 force`	Forcefully applies the changes to the enhanced zone database and closes the session created by another user.
	`switch(config)# no zone commit vsan 2`	Discards the changes to the enhanced zone database and closes the session.
	`switch(config)# no zone commit vsan 3 force`	Forcefully discards the changes to the enhanced zone database and closes the session created by another user.

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# config t
switch(config)# no zone commit vsan 2
```

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

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

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

Creating Attribute Groups

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

To configure attribute groups, follow these steps:

Step 1 Create an attribute group.

```
switch# conf t
switch(config)# zone-attribute-group name SampleAttributeGroup vsan 2
switch(config-attribute-group)#
```

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

```
switch(config-attribute-group)# readonly
switch(config-attribute-group)# broadcast
switch(config-attribute-group)# qos priority medium
```

Step 3 Attach the attribute-group to a zone.

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

Step 4 Activate the zone set.

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

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

Merging the Database

The merge behavior depends on the fabric-wide merge control setting:

- Restrict—If the two database are not identical, the ISLs between the switches are isolated.
- Allow—The two databases are merged using the merge rules specified in Table 23-3.

Table 23-3 Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name[1] but different zones, aliases, and attributes groups.		Successful.	The union of the local and adjacent databases.
The databases contains a zone, zone alias, or zone attribute group object with same name[1] but different members.		Failed.	ISLs are isolated.
Empty.	Contains data.	Successful.	The adjacent database information populates the local database.
Contains data.	Empty.	Successful.	The local database information populates the adjacent database.

1. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets.

Caution Remove all non-pWWN-type zone entries on all MDS switches running Cisco SAN-OS prior to merging fabrics if there is a Cisco MDS 9020 switch running FabricWare in the adjacent fabric.

The Merge Process

The merge process operates as follows:

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

Configuring Zone Merge Control Policies

To configure merge control policies, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# zone merge-control restrict vsan 4`	Configures a restricted merge control setting for this VSAN.
	`switch(config)# no zone merge-control restrict vsan 2`	Defaults to using the allow merge control setting for this VSAN.
	`switch(config)# zone commit vsan 4`	Commits the changes made to VSAN 4.

Default Zone Policies

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# zone default-zone permit vsan 5`	Permits traffic flow to default zone members.
	`switch(config)# no zone default-zone permit vsan 3`	Denies traffic flow to default zone members and reverts to factory default.
Step 3	`switch(config)# zone commit vsan 5`	Commits the changes made to VSAN 5.

Broadcasting a Zone

You can specify an enhanced zone to restrict broadcast frames generated by a member in this zone to members within that zone. Use this feature when the host or storage devices support broadcasting.

Table 23-4 identifies the rules for the delivery of broadcast frames.

Table 23-4 Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.

Tip If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# zone-attribute-group name BroadcastAttr vsan 2`	Configures the zone attribute group for the required VSAN.
	`switch(config)# no zone-attribute-group name BroadAttr vsan 1`	Removes the zone attribute group for the required VSAN.

Step 3	`switch(config-attribute-group)# broadcast`
	`switch(config-attribute-group)# exit`
	`switch(config)#`
		Creates a broadcast attribute for this group and exits this submode.

	`switch(config-attribute-group)# no broadcast`	Removes broadcast attribute for this group and exits this submode.
Step 4	`switch(config)# zone name BroadcastAttr vsan 2` `switch(config-zone)#`	Configures a zone named BroadcastAttr in VSAN 2.

Step 5	`switch(config-zone)# member pwwn 21:00:00:e0:8b:0b:66:56`
	`switch(config-zone)# member pwwn 21:01:00:e0:8b:2e:80:93`
	`switch(config-zone)# attribute-group name BroadcastAttr`
	`switch(config-zone)# exit`
	`switch(config)#`
		Adds the specified members to this zone and exits this submode.

Step 6	`switch(config)# zone commit vsan 1`
	`Commit operation initiated`
	`switch(config)# end`
		Applies the changes to the enhanced zone configuration and exits this submode.

Step 7	`switch# show zone vsan 1`
	`zone name BroadcastAttr vsan 1`
	`zone-attribute-group name BroadcastAttr vsan 1`
	`broadcast`
	`pwwn 21:00:00:e0:8b:0b:66:56`
	`pwwn 21:01:00:e0:8b:2e:80:93`
		Displays the broadcast configuration

Configuring System Default Zoning Settings

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system default zone default-zone permit`	Configures permit as the default zoning policy for new VSANs on the switch.
	`switch(config)# no system default zone default-zone permit`	Configures deny (default) as the default zoning policy for new VSANs on the switch.
Step 3	`switch(config)# system default zone distribute full`	Enables full zone database distribution as the default for new VSANs on the switch.
Step 4	`switch(config)# 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.

Note Since VSAN 1 is the default VSAN and is always present on the switch, the **system default zone** commands have no effect on VSAN 1.

Displaying Enhanced Zone Information

You can view any zone information by using the **show** command. See Examples 23-20 to 23-32.

Example 23-20 Displays the Active Zone Set Information for a Specified VSAN

```
switch# show zoneset active vsan 2
zoneset name testzoneset vsan 2
  zone name testzone vsan 2
    attribute read-only
    attribute broadcast
    attribute qos priority high
    pwwn 21:01:00:e0:8b:2e:a3:8a
    pwwn 22:00:00:0c:50:02:cb:59

  zone name $default_zone$ vsan 2
    attribute read-only
    attribute qos priority high
    attribute broadcast]
```

Example 23-21 Displays the Zone Set Information or a Specified VSAN

```
switch# show zoneset vsan 2
zoneset name testzoneset vsan 2
  zone name testzone vsan 2
    zone-attribute-group name testattgp vsan 2
      read-only
      broadcast
      qos priority high
    pwwn 21:01:00:e0:8b:2e:a3:8a
    pwwn 22:00:00:0c:50:02:cb:59

zoneset name testzoneset2 vsan 2
  zone name testzone2 vsan 2
    pwwn 21:01:00:e0:8b:2e:68:8a
    pwwn 22:00:00:0c:50:02:cb:80

zoneset name testzoneset3 vsan 2
  zone name testzone3 vsan 2
    pwwn 21:01:00:e0:8b:2e:68:8a
    pwwn 22:00:00:0c:50:02:cb:80
```

Example 23-22 Displays the Zone Attribute Group Information for a Specified VSAN

```
switch# show zone-attribute-group vsan 2
zone-attribute-group name $default_zone_attr_group$ vsan 2
  read-only
  qos priority high
  broadcast
zone-attribute-group name testattgp vsan 2
  read-only
  broadcast
  qos priority high
```

Example 23-23 Displays the fcalias Information for the Specified VSAN

```
switch# show fcalias vsan 2
fcalias name testfcalias vsan 2
 pwwn 21:00:00:20:37:39:b0:f4
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
```

Example 23-24 Displays the Zone Status for the Specified VSAN

```
switch# show zone status vsan 2
VSAN: 2 default-zone: permit distribute: active only Interop: 100
    mode:basic merge-control:allow session:none
    hard-zoning:enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:3  Zones:3 Aliases: 0 Attribute-groups: 2
Active Zoning Database :
    Name: testzoneset  Zonesets:1  Zones:2
Status:
```

Example 23-25 Displays an Active Zone Status for the Specified VSAN

```
switch# show zone status vsan 1
VSAN: 1 default-zone: permit distribute: full Interop: 100
    mode: enhanced merge-control: allow session: active <--------Indicates an active session.
    Hard zoning is enabled
Default zone:
    qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:4  Zones:4 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
    Database Not Available
Status: Set zoning mode complete at 10:36:48 Aug 18 2004
```

Example 23-26 Displays the Pending Zone Set Information for the VSAN to be Committed

```
switch# show zoneset pending vsan 2
No pending info found
```

Example 23-27 Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone pending vsan 2
No pending info found
```

Example 23-28 Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone-attribute-group pending vsan 2
No pending info found
```

Example 23-29 Displays the Pending Active Zone Set Information for the VSAN to be Committed

```
switch# show zoneset pending active vsan 2
No pending info found
```

Example 23-30 Displays the Difference Between the Pending and Effective Zone Information for the Specified VSAN

```
switch# show zone pending-diff vsan 2
zone name testzone vsan 2
   -  member pwwn 21:00:00:20:37:4b:00:a2
   +  member pwwn 21:00:00:20:37:60:43:0c
```

Exchange Switch Support (ESS) defines a mechanism for two switches to exchange various supported features (see Example 23-31).

Example 23-31 Displays the ESS Information for All Switches in the Specified VSAN

```
switch# show zone ess vsan 2
ESS info on VSAN 2 :
    Domain : 210, SWWN : 20:02:00:05:30:00:85:1f, Cap1 : 0xf3, Cap2 : 0x0
```

Example 23-32 Displays the Pending fcalias Information for the VSAN to be Committed

```
switch# show fcalias pending vsan 2
No pending info found
```

Compacting the Zone Database for Downgrading

Prior to Cisco SAN-OS Release 3.0(1), only 2000 zones are supported per VSAN. If you add more than 2000 zones to a VSAN, a configuration check is registered to indicate that downgrading to a previous release could cause you to lose the zones over the limit. To avoid the configuration check, delete the excess zones and compact the zone database for the VSAN. If there are 2000 zones or fewer after deleting the excess zones, the compacting process assigns new internal zone IDs and the configuration can be supported by Cisco SAN-OS Release 2.x or earlier. Perform this procedure for every VSAN on the switch with more than 2000 zones.

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

, To delete zones and compact the zone database for a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no zone name ExtraZone vsan 10`	Deletes a zone to reduce the number of zones to 2000 or fewer.
Step 3	`switch(config)# zone compact vsan 10`	Compacts the zone database for VSAN 10 to recover the zone ID released when a zone was deleted.

Zone and Zone Set Analysis

To better manage the zones and zone sets on your switch, you can display zone and zone set information using the **show zone analysis** command (see Example 23-33 through Example 23-35).

Example 23-33 Full Zoning Analysis

```
switch# show zone analysis vsan 1
Zoning database analysis vsan 1
 Full zoning database
   Last updated at: 15:57:10 IST Feb 20 2006
   Last updated by: Local [ CLI ]
   Num zonesets: 1
   Num zones: 1
   Num aliases: 0
   Num attribute groups: 0
   Formattted size: 36 bytes / 2048 Kb

 Unassigned Zones: 1
   zone name z1 vsan 1
```

Note The maximum size of the full zone database per VSAN is 2000 KB.

Example 23-34 Active Zoning Database Analysis

```
switch# show zone analysis active vsan 1
Zoning database analysis vsan 1
  Active zoneset: zs1 [*]
    Activated at: 08:03:35 UTC Nov 17 2005
    Activated by: Local [ GS ]
    Default zone policy: Deny
    Number of devices zoned in vsan: 0/2 (Unzoned: 2)
    Number of zone members resolved: 0/2 (Unresolved: 2)
    Num zones: 1
    Number of IVR zones: 0
    Number of IPS zones: 0
    Formattted size: 38 bytes / 2048 Kb
```

Note The maximum size of the active zone set database per VSAN is 2000 KB.

Example 23-35 Zone Set Analysis

```
switch# show zone analysis zoneset zs1 vsan 1
Zoning database analysis vsan 1
  Zoneset analysis: zs1
    Num zonesets: 1
    Num zones: 0
    Num aliases: 0
    Num attribute groups: 0
    Formattted size: 20 bytes / 2048 Kb
```

See the *Cisco MDS 9000 Family Command Reference* for the description of the information displayed in the command output.

Default Settings

Table 23-5 lists the default settings for basic zone parameters.

Table 23-5 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.
Zone based traffic priority	Low.
Read-only zones	Read-write attributes for all zones.
Broadcast frames	Sent to all Nx ports.
Broadcast zoning	Disabled.
Enhanced zoning	Disabled.

CHAPTER 24

Distributing Device Alias Services

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

This chapter includes the following sections:

About Device Aliases

When the port WWN of a device must be specified to configure different features (zoning, QoS, port security) in a Cisco MDS 9000 Family switch, you must assign the right 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 port WWN and use this name in all the configuration commands as required. These user-friendly names are referred to as *device aliases* in this chapter.

This section includes the following topics:

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of your 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 loosing data.

- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (see Chapter 6, "Using the CFS Infrastructure").
- When you configure zones, IVR zones, or QoS features using device aliases, and if you display these configurations, you see that the device aliases are displayed along with their respective pWWNs.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- The mapping between the pWWN and the device alias to which it is mapped must have a one-to-one relationship. A pWWN can be mapped to only one device alias and vice versa.
- 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
  - 1 to 9
  - - (hyphen) and _ (underscore)
  - $ (dollar sign) and ^ (up carat)

Zone Aliases Versus Device Aliases

Table 24-1 compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 24-1 Comparison Between Zone Aliases and Device Aliases

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

Device Alias Databases

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

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

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

This section includes the following topics:

About Device Alias Distribution

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

If you have not committed the changes and you enable distribution, then a commit task will fail.

Device Alias Statistics Cleanup

To clear device alias statistics (for debugging purposes), refer to the *Cisco MDS 9000 Family CLI Configuration Guide*.

Database Merge Guidelines

Refer to the "CFS Merge Support" section on page 6-8 for detailed concepts.

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two identical pWWNs are not mapped to two different device aliases.
- Verify that the combined number of the device aliases in both databases does not exceed 8191 (8K). For example, if database N has 6000 device aliases and database M has 2192 device aliases, this merge operation will fail.

Default Settings

Table 24-2 lists the default settings for device alias parameters.

Table 24-2 Default Device Alias Parameters

Parameters	Default
Database in use	Effective database.
Database to accept changes	Pending database.
Device alias fabric lock state	Locked with the first device alias task.

CHAPTER 25

Configuring Fibre Channel Routing Services and Protocols

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

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

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

About FSPF

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's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.

Note The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 25-1 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 25-1 Fault Tolerant Fabric

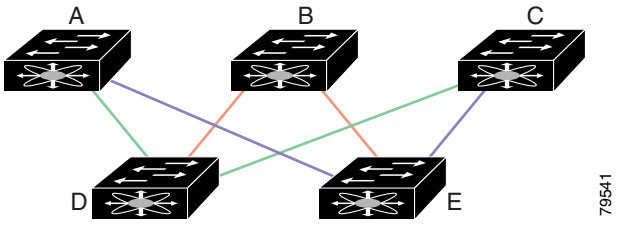

For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in Figure 25-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 25-2 shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

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

Figure 25-2 Fault Tolerant Fabric with Redundant Links

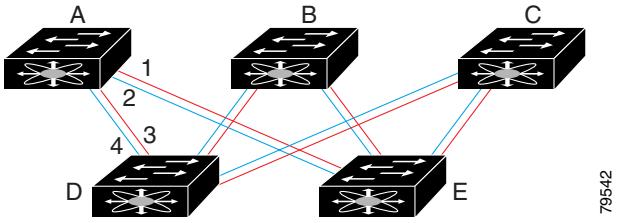

For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-Over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in Figure 25-3. Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100% utilization at 1 Gbps in two scenarios:

- Disabling the traffic link by physically removing the cable (see Table 25-1).
- Shutting down either switch 1 or switch 2 (see Table 25-2).

Figure 25-3 Fail-Over Scenario Using Traffic Generators

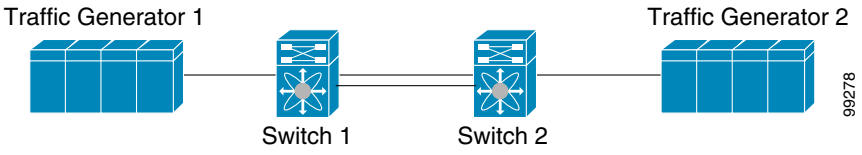

Table 25-1 Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	**Switch 2**	**Switch 1**	**Switch 2**
110 msec (~2K frame drops)		130+ msec (~4k frame drops)	
100 msec (hold time when a signal loss is reported as mandated by the standard)			

Table 25-2 Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	**Switch 2**	**Switch 1**	**Switch 2**
~0 msec (~8 frame drops)	110 msec (~2K frame drops)	130+ msec (~4K frame drops)	
No hold time needed	Signal loss on switch 1	No hold time needed	Signal loss on switch 1

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

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

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

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

This section includes the following topics:

About SPF Computational Hold Times

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

About Link State Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. Table 25-3 displays the default settings for switch responses.

Table 25-3 LSR Default Settings

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

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

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

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fspf config vsan 1`	Enters FSPF global configuration mode for the specified VSAN.
Step 3	`switch-config-(fspf-config)# spf static`	Forces static SPF computation for the dynamic (default) incremental VSAN.
Step 4	`switch-config-(fspf-config)# spf hold-time 10`	Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0. **Note** If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.
Step 5	`switch-config-(fspf-config)# region 7`	Configures the autonomous region for this VSAN and specifies the region ID (7).

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no fspf config vsan 3`	Deletes the FSPF configuration for VSAN 3.

Enabling or Disabling FSPF

To enable or disable FSPF routing protocols, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fspf enable vsan 7`	Enables the FSPF routing protocol in VSAN 7.
	`switch(config)# no fspf enable vsan 5`	Disables the FSPF routing protocol in VSAN 5.

Clearing FSPF Counters for the VSAN

To clear the FSPF statistics counters for the entire VSAN, follow this step:

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

FSPF Interface Configuration

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

This section includes the following topics:

About FSPF Link Cost

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

Configuring FSPF Link Cost

To configure FSPF link cost, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/4` `switch(config-if)#`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	`switch(config-if)# fspf cost 5 vsan 90`	Configures the cost for the selected interface in VSAN 90.

About Hello Time Intervals

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

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

Configuring Hello Time Intervals

To configure the FSPF Hello time interval, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/4` `switch(config-if)#`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	`switch(config-if)# fspf hello-interval 15 vsan 175` `switch(config-if)#`	Specifies the hello message interval (15 seconds) to verify the health of the link in VSAN 175. The default is 20 seconds.

About Dead Time Intervals

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

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

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

Configuring Dead Time Intervals

To configure the FSPF dead time interval, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/4` `switch(config-if)#`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	`switch(config-if)# fspf dead-interval 25 vsan 7` `switch(config-if)#`	Specifies the maximum interval for VSAN 7 before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

About Retransmitting Intervals

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

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

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/4` `switch(config-if)#`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	`switch(config-if)# fspf retransmit-interval 15 vsan 12` `switch(config-if)#`	Specifies the retransmit time interval for unacknowledged link state updates in VSAN 12. The default is 5 seconds.

About Disabling FSPF for Specific Interfaces

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

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

Disabling FSPF for Specific Interfaces

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

To disable FSPF for a specific interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/4` `switch(config-if)#`	Configures a specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	`switch(config-if)# fspf passive vsan 1` `switch(config-if)#`	Disables the FSPF protocol for the specified interface in the specified VSAN.
	`switch(config-if)# no fspf passive vsan 1` `switch(config-if)#`	Reenables the FSPF protocol for the specified interface in the specified VSAN.

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

Clearing FSPF Counters for an Interface

To clear the FSPF statistics counters for an interface, follow this step:

	Command	Purpose
Step 4	`switch# clear fspf counters vsan 200 interface fc1/1`	Clears the FSPF statistics counters for the specified interface in VSAN 200.

FSPF Routes

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

This section includes the following topics:

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see Figure 25-4).

Figure 25-4 Fibre Channel Routes

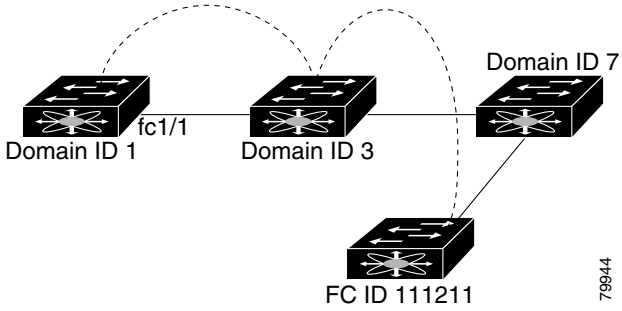

Note Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

Configuring Fibre Channel Routes

To configure a Fibre Channel route, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# fcroute 0x111211 interface fc1/1 domain 3 vsan 2` `switch(config)#`	Configures the route for the specified Fibre Channel interface and domain. In this example, interface fc1/1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch.
	`switch(config)# fcroute 0x111211 interface port-channel 1 domain 3 vsan 4` `switch(config)#`	Configures the route for the specified PortChannel interface and domain. In this example, interface port-channel 1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch.
	`switch(config)# fcroute 0x031211 interface fc1/1 domain 3 metric 1 vsan 1` `switch(config-if)#`	Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route. If the remote destination option is not specified, the default is direct.
	`switch(config)# fcroute 0x111112 interface fc1/1 domain 3 metric 3 remote vsan 3`	Adds a static route to the RIB. If this is an active route and the FIB[1] records are free, it is also added to the FIB. If the cost (metric) of the route is not specified, the default is 10.
Step 3	`switch(config)# fcroute 0x610000 0xff0000 interface fc 1/1 domain 1 vsan 2` `switch(config)#`	Configures the netmask for the specified route in interface fc1/1 (or PortChannel). You can specify one of three routes: 0xff0000 matches only the domain, 0xffff00 matches the domain and the area, 0xffffff matches the domain, area, and port.

1. FIB = Forwarding Information Base

About Broadcast and Multicast Routing

Broadcast and multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive a loop-free distribution tree for multicast and broadcast routing in a VSAN.

Caution All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

To interoperate with other vendor switches (following FC-SW3 guidelines), the SAN-OS software uses the lowest domain switch as the root to compute the multicast tree in interop mode.

About Multicast Root Switch

By default, the **native** (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could face potential loop and frame-drop problems.

Note The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

Use the **mcast root lowest vsan** command to change the multicast root from the principal switch to lowest domain switch.

Setting the Multicast Root Switch

To use the lowest domain switch for the multicast tree computation, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# mcast root lowest vsan 1`	Uses the lowest domain switch to compute the multicast tree.
	`switch(config)# mcast root principal vsan 1`	Defaults to using the principal switch to compute the multicast tree.

To display the configured and operational multicast mode and the selected root domain, use the **show mcast** command.

```
switch# show mcast vsan 1
Multicast root for VSAN 1
       Configured root mode : Principal switch
       Operational root mode : Principal switch
       Root Domain ID : 0xef(239)
```

In-Order Delivery

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

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

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

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

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

This section includes the following topics:

About Reordering Network Frames

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

Figure 25-5 Route Change Delivery

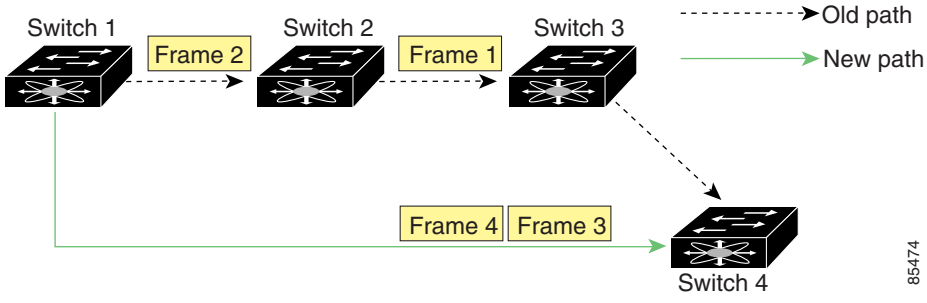

In Figure 25-5, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

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

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

About Reordering PortChannel Frames

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

Figure 25-6 Link Congestion Delivery

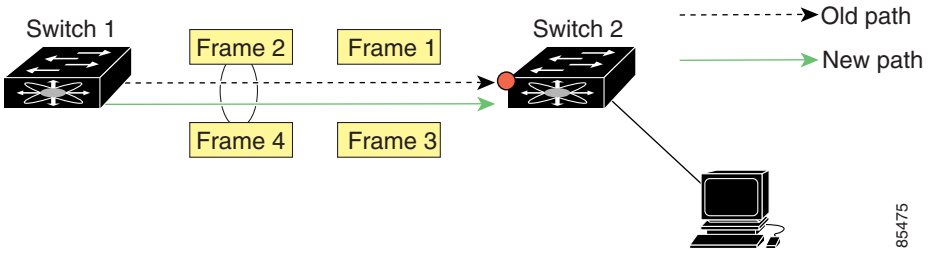

In Figure 25-6, the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

The the in-order delivery feature attempts to minimize the number of frames dropped during PortChannel link changes when the in-order delivery is enabled by sending a request to the remote switch on the PortChannel to flush all frames for this PortChannel.

Note Both switches on the PortChannel must be running Cisco SAN-OS Release 3.0(1) for this IOD enhancement. For earlier releases, IOD waits for the switch latency period before sending new frames.

When the in-order delivery guarantee feature is enabled and a PortChannel link change occurs, the frames crossing the PortChannel are treated as follows:

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

Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped. See the "Configuring the Drop Latency Time" section on page 25-17.

About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.

Tip We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

Enabling In-Order Delivery Globally

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

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

Note Enable in-order delivery on the entire switch before performing a downgrade to Cisco MDS SAN-OS Release1.3(3) or earlier.

To enable in-order delivery for the switch, follow these steps:

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

Enabling In-Order Delivery for a VSAN

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

To use the lowest domain switch for the multicast tree computation, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# in-order-guarantee vsan 3452`	Enables in-order delivery in VSAN 3452.
	`switch(config)# no in-order-guarantee vsan 101`	Reverts the switch to the factory defaults and disables the in-order delivery feature in VSAN 101.

Displaying the In-Order Delivery Status

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

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed

VSAN specific settings
vsan 1 inorder delivery:guaranteed
vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
```

```
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed
```

Configuring the Drop Latency Time

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

To configure the network and the switch drop latency time, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcdroplatency network 5000`	Configures network drop latency time to be 5000 msec for the network. The valid range is 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
	`switch(config)# fcdroplatency network 6000 vsan 3`	Configures network drop latency time to be 6000 msec for VSAN 3.
	`switch(config)# no fcdroplatency network 4500`	Removes the current fcdroplatecy network configuration (4500) and reverts the switch to the factory defaults.
Step 3	`switch(config)# fcdroplatency switch 4000`	Configures switch drop latency time to be 4000 msec for the switch. The valid range is 0 to 60000 msec. The default is 500 msec. **Note** The switch drop latency parameter should have the same value in all the switches in the network
	`switch(config)# no fcdroplatency switch 4500`	Removes the current fcdroplatecy switch configuration (4500) and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplatency** command (see Example 25-1).

Example 25-1 Displays Administrative Distance

```
switch# show fcdroplatency
switch latency value:500 milliseconds
global network latency value:2000 milliseconds

VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
```

```
vsan 460 network latency:500 milliseconds
```

Flow Statistics Configuration

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

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

This section includes the following topics:

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics for Generation 1 modules, and 2 K entries for Generation 2 modules. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Generation 1 modules allow a maximum of 1024 flow statements per module. Generation 2 modules allow a maximum of 2048-128 flow statements per module.

Counting Aggregated Flow Statistics

To count the aggregated flow statistics for a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcflow stats aggregated module 1 index 1005 vsan 1` `switch(config)#`	Enables the aggregated flow counter.
	`switch(config)# no fcflow stats aggregated module 1 index 1005 vsan 1` `switch(config)#`	Disables the aggregated flow counter.

Counting Individual Flow Statistics

To count the flow statistics for a source and destination FC ID in a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcflow stats module 1 index 1 0x145601 0x5601ff 0xffffff vsan 1` `switch(config)#`	Enables the flow counter. **Note** The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xffffff.
	`switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2` `switch(config)#`	Disables the flow counter.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter (see Examples 25-2 and 25-3).

Example 25-2 Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Example 25-3 Clears Flow Counters for Source and Destination FC IDs

```
switch# clear fcflow stats module 2 index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics (see Example 25-4 to 25-6).

Example 25-4 Displays Aggregated Flow Details for the Specified Module

```
switch# show fcflow stats aggregated module 2
Idx  VSAN # frames # bytes
---- ---- -------- -------
0000 4    387,653  674,235,875
0001 6    34,402   2,896,628
```

Example 25-5 Displays Flow Details for the Specified Module

```
switch# show fcflow stats module 2
Idx  VSAN D ID        S ID        mask      # frames # bytes
---- ---- ----------- ----------- -----     -------- -------
0000 4    032.001.002 007.081.012 ff.ff.ff   387,653  674,235,875
0001 6    004.002.001 019.002.004 ff.00.00   34,402   2,896,628
```

Example 25-6 Displays Flow Index Usage for the Specified Module

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

Displaying Global FSPF Information

Example 25-7 displays global FSPF information for a specific VSAN:

- Domain number of the switch.
- Autonomous region for the switch.
- Min_LS_arrival: minimum time that must elapse before the switch accepts LSR updates.
- Min_LS_interval: minimum time that must elapse before the switch can transmit an LSR.

Tip If the Min_LS_interval is higher than 10 seconds, the graceful shutdown feature is not implemented.

- LS_refresh_time: interval time lapse between refresh LSR transmissions.
- Max_age: maximum time aa LSR can stay before being deleted.

Example 25-7 Displays FSPF Information for a Specified VSAN

```
switch# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x65(101)
Number of LSRs = 3, Total Checksum = 0x0001288b

Protocol constants :
   LS_REFRESH_TIME = 1800 sec
   MAX_AGE         = 3600 sec

Statistics counters :
   Number of LSR that reached MaxAge = 0
   Number of SPF computations        = 7
   Number of Checksum Errors         = 0
   Number of Transmitted packets :  LSU 65 LSA 55 Hello 474 Retranmsitted LSU 0
   Number of received packets :  LSU 55 LSA 60 Hello 464 Error packets 10
```

Displaying the FSPF Database

Example 25-8 displays a summary of the FSPF database for a specified VSAN. If other parameters are not specified, all LSRs in the database are displayed:

- LSR type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

You could narrow the display to obtain specific information by issuing additional parameters for the domain ID of the LSR owner. For each interface, the following information is also available:

- Domain ID of the neighboring switch
- E port index
- Port index of the neighboring switch
- Link type and cost

Example 25-8 Displays FSPF Database Information

```
switch# show fspf database vsan 1

FSPF Link State Database for VSAN 1 Domain 0x0c(12)
LSR Type                = 1
Advertising domain ID   = 0x0c(12)
LSR Age                 = 1686
LSR Incarnation number  = 0x80000024
LSR Checksum            = 0x3caf
Number of links         = 2
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
--------------------------------------------------------------------------------
   0x65(101) 0x0000100e     0x00001081               1          500
   0x65(101) 0x0000100f     0x00001080               1          500

FSPF Link State Database for VSAN 1 Domain 0x65(101)
LSR Type                = 1
Advertising domain ID   = 0x65(101)
LSR Age                 = 1685
LSR Incarnation number  = 0x80000028
LSR Checksum            = 0x8443
Number of links         = 6
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
--------------------------------------------------------------------------------
   0xc3(195) 0x00001085     0x00001095               1          500
   0xc3(195) 0x00001086     0x00001096               1          500
   0xc3(195) 0x00001087     0x00001097               1          500
   0xc3(195) 0x00001084     0x00001094               1          500
    0x0c(12) 0x00001081     0x0000100e               1          500
    0x0c(12) 0x00001080     0x0000100f               1          500

FSPF Link State Database for VSAN 1 Domain 0xc3(195)
LSR Type                = 1
Advertising domain ID   = 0xc3(195)
LSR Age                 = 1686
LSR Incarnation number  = 0x80000033
```

```
LSR Checksum            = 0x6799
Number of links         = 4
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
-------------------------------------------------------------------------------
   0x65(101) 0x00001095     0x00001085               1          500
   0x65(101) 0x00001096     0x00001086               1          500
   0x65(101) 0x00001097     0x00001087               1          500
   0x65(101) 0x00001094     0x00001084               1          500
```

Displaying FSPF Interfaces

Example 25-9 displays the following information for each selected interface.

- Link cost
- Timer values
- Neighbor's domain ID (if known)
- Local interface number
- Remote interface number (if known)
- FSPF state of the interface
- Interface counters

Example 25-9 Displays FSPF Interface Information

```
switch# show fspf vsan 1 interface fc1/1
FSPF interface fc1/1 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x0c(12), Neighbor Interface index is 0x0f100000
Statistics counters :
   Number of packets received : LSU  8  LSA  8  Hello 118  Error packets 0
   Number of packets transmitted : LSU  8  LSA  8  Hello 119  Retransmitted LSU 0
   Number of times inactivity timer expired for the interface = 0
```

Default Settings

Table 25-4 lists the default settings for FSPF features.

Table 25-4 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.

Table 25-4 Default FSPF Settings (continued)

Parameters	Default
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
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 26

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login database, the name server features, the Fabric-Device Management Interface, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family. It includes the following sections:

FLOGI

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.

Displaying FLOGI Details

See Examples 26-1 to 26-4.

Example 26-1 Displays Details on the FLOGI Database

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
sup-fc0    2     0xb30100  10:00:00:05:30:00:49:63  20:00:00:05:30:00:49:5e
fc9/13     1     0xb200e2  21:00:00:04:cf:27:25:2c  20:00:00:04:cf:27:25:2c
fc9/13     1     0xb200e1  21:00:00:04:cf:4c:18:61  20:00:00:04:cf:4c:18:61
fc9/13     1     0xb200d1  21:00:00:04:cf:4c:18:64  20:00:00:04:cf:4c:18:64
fc9/13     1     0xb200ce  21:00:00:04:cf:4c:16:fb  20:00:00:04:cf:4c:16:fb
```

```
fc9/13      1      0xb200cd  21:00:00:04:cf:4c:18:f7  20:00:00:04:cf:4c:18:f7

Total number of flogi = 6.
```

Example 26-2 Displays the FLOGI Database by Interface

```
switch# show flogi database interface fc1/11
INTERFACE   VSAN    FCID        PORT NAME                NODE NAME
----------  ------  ---------   ----------------------   ---------------------
fc1/11      1       0xa002ef    21:00:00:20:37:18:17:d2  20:00:00:20:37:18:17:d2
fc1/11      1       0xa002e8    21:00:00:20:37:38:a7:c1  20:00:00:20:37:38:a7:c1
fc1/11      1       0xa002e4    21:00:00:20:37:6b:d7:18  20:00:00:20:37:6b:d7:18
fc1/11      1       0xa002e2    21:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45
fc1/11      1       0xa002e1    21:00:00:20:37:39:90:6a  20:00:00:20:37:39:90:6a
fc1/11      1       0xa002e0    21:00:00:20:37:36:0b:4d  20:00:00:20:37:36:0b:4d
fc1/11      1       0xa002dc    21:00:00:20:37:5a:5b:27  20:00:00:20:37:5a:5b:27
fc1/11      1       0xa002da    21:00:00:20:37:18:6f:90  20:00:00:20:37:18:6f:90
fc1/11      1       0xa002d9    21:00:00:20:37:5b:cf:b9  20:00:00:20:37:5b:cf:b9
fc1/11      1       0xa002d6    21:00:00:20:37:46:78:97  0:00:00:20:37:46:78:97

Total number of flogi = 10.
```

Example 26-3 Displays the FLOGI Database by VSAN

```
switch# show flogi database vsan 1
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/3      1     0xef02ef  22:00:00:20:37:18:17:d2  20:00:00:20:37:18:17:d2
fc1/3      1     0xef02e8  22:00:00:20:37:38:a7:c1  20:00:00:20:37:38:a7:c1
fc1/3      1     0xef02e4  22:00:00:20:37:6b:d7:18  20:00:00:20:37:6b:d7:18
fc1/3      1     0xef02e2  22:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45
fc1/3      1     0xef02e1  22:00:00:20:37:39:90:6a  20:00:00:20:37:39:90:6a
fc1/3      1     0xef02e0  22:00:00:20:37:36:0b:4d  20:00:00:20:37:36:0b:4d
fc1/3      1     0xef02dc  22:00:00:20:37:5a:5b:27  20:00:00:20:37:5a:5b:27
fc1/3      1     0xef02da  22:00:00:20:37:18:6f:90  20:00:00:20:37:18:6f:90
fc1/3      1     0xef02d9  22:00:00:20:37:5b:cf:b9  20:00:00:20:37:5b:cf:b9
fc1/3      1     0xef02d6  22:00:00:20:37:46:78:97  20:00:00:20:37:46:78:97

Total number of flogi = 10.
```

Example 26-4 Displays the FLOGI Database by FC ID

```
switch# show flogi database fcid 0xef02e2
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/3      1     0xef02e2  22:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45

Total number of flogi = 1.
```

See the "Default Company ID list" section on page 29-9 and the "Loop Monitoring" section on page 58-15.

Name Server Proxy

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

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

This section includes the following topics:

About Registering Name Server Proxies

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

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

Registering Name Server Proxies

To register the name server proxy, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcns proxy-port 21:00:00:e0:8b:00:26:d0 vsan 2`	Configures a proxy port for the specified VSAN.

About Rejecting Duplicate pWWN

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

Rejecting Duplicate pWWNs

To reject duplicate pWWNs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcns reject-duplicate-pwwn vsan 1`	Logs out devices when they log into the fabric if the pWWNs already exist.
	`switch(config)# no fcns reject-duplicate-pwwn vsan 1`	Overwrites the first device's entry in the name server database with the new device having the same pwwn (default).

About Name Server Database Entries

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

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

Displaying Name Server Database Entries

Use the **show fcns** command to display the name server database and statistical information for a specified VSAN or for all VSANs (see Examples 26-5 to 26-8).

Example 26-5 Displays the Name Server Database

```
switch# show fcns database
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000    N     50:06:0b:00:00:10:a7:80                   scsi-fcp fc-gs
0x010001    N     10:00:00:05:30:00:24:63 (Cisco)           ipfc
0x010002    N     50:06:04:82:c3:a0:98:52 (Company 1)       scsi-fcp 250
0x010100    N     21:00:00:e0:8b:02:99:36 (Company A)       scsi-fcp
0x020000    N     21:00:00:e0:8b:08:4b:20 (Company A)
0x020100    N     10:00:00:05:30:00:24:23 (Cisco)        ipfc
0x020200    N     21:01:00:e0:8b:22:99:36 (Company A)       scsi-fcp
```

Example 26-6 Displays the Name Server Database for the Specified VSAN

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

```
Total number of entries = 4
```

Example 26-7 Displays the Name Server Database Details

```
switch# show fcns database detail
------------------------
VSAN:1     FCID:0x030001
------------------------
port-wwn (vendor)     :10:00:00:05:30:00:25:a3 (Cisco)
node-wwn              :20:00:00:05:30:00:25:9e
class                 :2,3
node-ip-addr          :0.0.0.0
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:ipfc
symbolic-port-name    :
symbolic-node-name    :
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :00:00:00:00:00:00:00:00
hard-addr             :0x000000
------------------------
VSAN:1     FCID:0xec0200
------------------------
port-wwn (vendor)     :10:00:00:5a:c9:28:c7:01
node-wwn              :10:00:00:5a:c9:28:c7:01
class                 :3
node-ip-addr          :0.0.0.0
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:
symbolic-port-name    :
symbolic-node-name    :
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :22:0a:00:05:30:00:26:1e
hard-addr             :0x000000
Total number of entries = 2
```

Example 26-8 Displays the Name Server Statistics

```
switch# show fcns statistics
registration requests received = 27
deregistration requests received = 0
queries received = 57
queries sent = 10
reject responses sent = 14
RSCNs received = 0
RSCNs sent = 0
```

FDMI

Cisco MDS 9000 Family switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel Host Bus Adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the SAN-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

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

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

Displaying FDMI

Use the **show fdmi** command to display the FDMI database information (see Examples 26-9 to 26-11).

Example 26-9 Displays All HBA Management Servers

```
switch# show fdmi database
Registered HBA List for VSAN 1
  10:00:00:00:c9:32:8d:77
  21:01:00:e0:8b:2a:f6:54
switch# show fdmi database detail
Registered HBA List for VSAN 1
--------------------------------
HBA-ID: 10:00:00:00:c9:32:8d:77
--------------------------------
Node Name        :20:00:00:00:c9:32:8d:77
Manufacturer     :Emulex Corporation
Serial Num       :0000c9328d77
Model            :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver     :2002606D
Driver Ver       :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver          :3.11A0
Firmware Ver     :3.90A7
OS Name/Ver      :Window 2000
CT Payload Len   :1300000
  Port-id: 10:00:00:00:c9:32:8d:77
--------------------------------
HBA-ID: 21:01:00:e0:8b:2a:f6:54
--------------------------------
Node Name        :20:01:00:e0:8b:2a:f6:54
Manufacturer     :QLogic Corporation
Serial Num       :\74262
Model            :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver     :FC5010409-10
Driver Ver       :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver          :1.24
Firmware Ver     :03.02.13.
OS Name/Ver      :500
CT Payload Len   :2040
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

Example 26-10 Displays HBA Details for a Specified VSAN

```
switch# show fdmi database detail vsan 1
Registered HBA List for VSAN 1
```

```
------------------------------
HBA-ID: 10:00:00:00:c9:32:8d:77
------------------------------
Node Name        :20:00:00:00:c9:32:8d:77
Manufacturer     :Emulex Corporation
Serial Num       :0000c9328d77
Model            :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver     :2002606D
Driver Ver       :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver          :3.11A0
Firmware Ver     :3.90A7
OS Name/Ver      :Window 2000
CT Payload Len   :1300000
  Port-id: 10:00:00:00:c9:32:8d:77
------------------------------
HBA-ID: 21:01:00:e0:8b:2a:f6:54
------------------------------
Node Name        :20:01:00:e0:8b:2a:f6:54
Manufacturer     :QLogic Corporation
Serial Num       :\74262
Model            :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver     :FC5010409-10
Driver Ver       :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver          :1.24
Firmware Ver     :03.02.13.
OS Name/Ver      :500
CT Payload Len   :2040
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

Example 26-11 Displays Details for the Specified HBA Entry

```
switch# show fdmi database detail hba-id 21:01:00:e0:8b:2a:f6:54 vsan 1

Node Name        :20:01:00:e0:8b:2a:f6:54
Manufacturer     :QLogic Corporation
Serial Num       :\74262
Model            :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver     :FC5010409-10
Driver Ver       :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver          :1.24
Firmware Ver     :03.02.13.
OS Name/Ver      :500
CT Payload Len   :2040
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

RSCN

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- 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.

This section includes the following topics:

About RSCN Information

Apart from sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.

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

Displaying RSCN Information

Use the **show rscn** command to display RSCN information (see Examples 26-12 and 26-13).

Example 26-12 Displays Register Device Information

```
switch# show rscn scr-table vsan 1
SCR table for VSAN: 1
---------------------------------------------
FC-ID         REGISTERED FOR
---------------------------------------------
0x1b0300      fabric detected rscns
Total number of entries = 1
```

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

Example 26-13 Displays RSCN Counter Information

```
switch# show rscn statistics vsan 1
Statistics for VSAN: 1
-------------------------
Number of SCR received          = 8
Number of SCR ACC sent          = 8
Number of SCR RJT sent          = 0
```

```
Number of RSCN received         = 0
Number of RSCN sent             = 24
Number of RSCN ACC received     = 24
Number of RSCN ACC sent         = 0
Number of RSCN RJT received     = 0
Number of RSCN RJT sent         = 0
Number of SW-RSCN received      = 6
Number of SW-RSCN sent          = 15
Number of SW-RSCN ACC received = 15
Number of SW-RSCN ACC sent      = 6
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent      = 0
```

About the multi-pid Option

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

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

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

Configuring the multi-pid Option

To configure the **multi-pid** option, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn multi-pid vsan 105`	Sends RSCNs in a multi-pid format for VSAN 105.

Suppressing Domain Format SW-RSCNs

A domain format SW-RSCN is sent whenever the local switch name or the local switch management IP address changes. This SW-RSCN is sent to all other domains and switches over the ISLs. The remote switches can issue GMAL and GIELN commands to the switch that initiated the domain format SW-RSCN to determine what changed. Domain format SW-RSCNs can cause problems with some non-Cisco MDS switches (refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*).

To suppress the transmission of these SW RSCNs over an ISL, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn suppress domain-swrscn vsan 105`	Suppresses transmission of domain format SW-RSCNs for VSAN 105.

Note You cannot suppress transmission of port address or area address format RSCNs.

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 (like ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

Use the **clear rscn statistics** command to clear the RSCN statistics for the specified VSAN.

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

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

```
switch# show rscn statistics vsan 1
Statistics for VSAN: 1
-----------------------
Number of SCR received         = 0
Number of SCR ACC sent         = 0
Number of SCR RJT sent         = 0
Number of RSCN received        = 0
Number of RSCN sent            = 0
Number of RSCN ACC received    = 0
Number of RSCN ACC sent        = 0
Number of RSCN RJT received    = 0
Number of RSCN RJT sent        = 0
Number of SW-RSCN received     = 0
Number of SW-RSCN sent         = 0
Number of SW-RSCN ACC received = 0
Number of SW-RSCN ACC sent     = 0
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent     = 0
```

Configuring the RSCN Timer

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

Note The RSCN timer value must be the same on all switches in the VSAN. See the "RSCN Timer Configuration Distribution" section on page 26-11.

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

To configure the RSCN timer, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn distribute`	Enables RSCN timer configuration distribution.
Step 3	`switch(config)# rscn event-tov 300 vsan 10`	Sets the event time-out value in milliseconds for the selected VSAN. In this example the event time-out value is set to 300 milliseconds for VSAN 12. The range is 0 to 2000 milliseconds. Setting a zero (0) value disables the timer.
	`switch(config)# no rscn event-tov 300 vsan 10`	Reverts to the default value (2000 milliseconds for Fibre Channel VSANs or 1000 milliseconds for FICON VSANs).
Step 4	`switch(config)# rscn commit vsan 10`	Commits the RSCN timer configuration to be distributed to the switches in VSAN 10.

Verifying the RSCN Timer Configuration

You verify the RSCN timer configuration using the **show rscn event-tov vsan** command.

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

RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs. See Chapter 6, "Using the CFS Infrastructure."

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

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

Note Only the RSCN timer configuration is distributed.

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

This section includes the following topics:

Note You can determine the compatibility when downgrading to an earlier Cisco MDS SAN-OS release using **show incompatibility system** command. You must disable RSCN timer distribution support before downgrading to an earlier release.

Note By default, the RSCN timer distribution capability is disabled and is compatible when upgrading from any Cisco MDS SAN-OS release earlier to 3.0.

Note For CFS distribution to operate correctly for the RSCN timer configuration, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later.

Enabling RSCN Timer Configuration Distribution

To enable RSCN timer configuration distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn distribute`	Enables RSCN timer distribution.
	`switch(config)# no rscn distribute`	Disables (default) RSCN timer distribution.

Locking the Fabric

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

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

Committing the RSCN Timer Configuration Changes

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

To commit RSCN timer configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn commit vsan 10`	Commits the RSCN timer changes.

Discarding the RSCN Timer Configuration Changes

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

To discard RSCN timer configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rscn abort vsan 10`	Discards the RSCN timer changes and clears the pending configuration database.

Clearing a Locked Session

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

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

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

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

Displaying RSCN Configuration Distribution Information

Use the **show cfs application name rscn** command to display the registration status for RSCN configuration distribution.

```
switch# show cfs application name rscn

 Enabled         : Yes
 Timeout         : 5s
 Merge Capable   : Yes
 Scope           : Logical
```

Use the **show rscn session status vsan** command to display session status information for RSCN configuration distribution.

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

```
switch# show rscn session status vsan 1
Session Parameters for VSAN: 1
------------------------------
Last Action                : Commit
Last Action Result         : Success
Last Action Failure Reason : None
```

Use the **show rscn pending** command to display the set of configuration commands that would take effect when you commit the configuration.

Note The pending database includes both existing and modified configuration.

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

Use the **show rscn pending-diff** command to display the difference between pending and active configurations. The following example shows the time-out value for VSAN 10 was changed from 2000 milliseconds (default) to 300 milliseconds.

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

Default Settings

Table 26-1 lists the default settings for RSCN.

Table 26-1 Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs 1000 milliseconds for FICON VSANs
RSCN timer configuration distribution	Disabled

CHAPTER 27

Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

About SCSI LUN Discovery

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

The name server requires LUN information for the following reasons:

- To display LUN storage device information so an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

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

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

This section includes the following topics:

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

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

Starting SCSI LUN Discovery

To start SCSI LUN discovery, follow this step:

	Command	Purpose
Step 1	`switch# discover scsi-target local os all` `discovery started`	Discovers local SCSI targets for all operating systems (OS). The operating system options are **aix**, **all**, **hpux**, **linux**, **solaris**, or **windows**
	`switch# discover scsi-target remote os aix` `discovery started`	Discovers remote SCSI targets assigned to the AIX OS.

	`switch# discover scsi-target vsan 1 fcid 0x9c03d6`
	`discover scsi-target vsan 1 fcid 0x9c03d6`
	`VSAN: 1 FCID: 0x9c03d6 PWWN:`
	`00:00:00:00:00:00:00:00`
	`PRLI RSP: 0x01 SPARM: 0x0012`
	`SCSI TYPE: 0 NLUNS: 1`
	`Vendor: Company 4 Model: ST318203FC Rev: 0004`
	`Other: 00:00:02:32:8b:00:50:0a`
		Discovers SCSI targets for the specified VSAN (1) and FC ID (0x9c03d6).

	`switch# discover scsi-target custom-list os linux` `discovery started`	Discovers SCSI targets from the customized list assigned to the Linux OS.

About Initiating Customized Discovery

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

Initiating Customized Discovery

To initiate a customized discovery, follow this step:

	Command	Purpose
Step 1	`switch# discover custom-list add vsan 1 domain 0X123456`	Adds the specified entry to the custom list.
	`switch# discover custom-list delete vsan 1 domain 0X123456`	Deletes the specified domain ID from the custom list.

Displaying SCSI LUN Information

Use the **show scsi-target** and **show fcns database** commands to display the results of the discovery. See Examples 27-1 to 27-8.

Example 27-1 Displays the Discovered Targets

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

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

Example 27-2 Displays the FCNS Database

```
switch# show fcns database

VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xeb0000    N     21:01:00:e0:8b:2a:f6:54 (Qlogic)        scsi-fcp:init
0xeb0201    NL    10:00:00:00:c9:32:8d:76 (Emulex)        scsi-fcp:init

Total number of entries = 2

VSAN 7:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xed0001    NL    21:00:00:04:cf:fb:42:f8 (Seagate)       scsi-fcp:target

Total number of entries = 1

VSAN 2002:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xcafe00    N     20:03:00:05:30:00:2a:20 (Cisco)         FICON:CUP

Total number of entries = 1
```

Example 27-3 Displays the Discovered Target Disks

```
switch# show scsi-target disk
-------------------------------------------------------------------------------
VSAN    FCID        PWWN                      VENDOR    MODEL              REV
-------------------------------------------------------------------------------
1       0x9c03d6    21:00:00:20:37:46:78:97   Company 4 ST318203FC         0004
1       0x9c03d9    21:00:00:20:37:5b:cf:b9   Company 4 ST318203FC         0004
1       0x9c03da    21:00:00:20:37:18:6f:90   Company 4 ST318203FC         0004
1       0x9c03dc    21:00:00:20:37:5a:5b:27   Company 4 ST318203FC         0004
1       0x9c03e0    21:00:00:20:37:36:0b:4d   Company 4 ST318203FC         0004
1       0x9c03e1    21:00:00:20:37:39:90:6a   Company 4 ST318203 CLAR18    3844
1       0x9c03e2    21:00:00:20:37:18:d2:45   Company 4 ST318203 CLAR18    3844
1       0x9c03e4    21:00:00:20:37:6b:d7:18   Company 4 ST318203 CLAR18    3844
1       0x9c03e8    21:00:00:20:37:38:a7:c1   Company 4 ST318203FC         0004
1       0x9c03ef    21:00:00:20:37:18:17:d2   Company 4 ST318203FC         0004
```

Example 27-4 Displays the Discovered LUNs for All Operating Systems

```
switch# show scsi-target lun os all
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-------------------------------------------------------------------------------
OS  LUN    Capacity Status  Serial Number    Device-Id
           (MB)
-------------------------------------------------------------------------------
WIN 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
AIX 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
SOL 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
LIN 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
HP  0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

Example 27-5 Displays the Discovered LUNs for the Solaris OS

```
switch# show scsi-target lun os solaris
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-------------------------------------------------------------------------------
OS  LUN    Capacity Status  Serial Number    Device-Id
           (MB)
-------------------------------------------------------------------------------
SOL 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

The following command displays the port WWN that is assigned to each OS (Windows, AIX, Solaris, Linux, or HPUX)

Example 27-6 Displays the pWWNs for each OS

```
switch# show scsi-target pwwn
------------------------------
OS      PWWN
------------------------------
WIN     24:91:00:05:30:00:2a:1e
AIX     24:92:00:05:30:00:2a:1e
SOL     24:93:00:05:30:00:2a:1e
LIN     24:94:00:05:30:00:2a:1e
HP      24:95:00:05:30:00:2a:1e
```

Example 27-7 Displays Customized Discovered Targets

```
switch# show scsi-target custom-list
---------------
VSAN    DOMAIN
---------------
1       56
```

Use the **show scsi-target auto-poll** command to verify automatic discovery of SCSI targets that come online. The internal uuid number indicates that a CSM or an IPS module is in the chassis.

Example 27-8 Displays Automatically Discovered Targets

```
switch# show scsi-target auto-poll
auto-polling is enabled, poll_start:0 poll_count:1 poll_type:0
USERS OF AUTO POLLING
---------------------
uuid:54
```

CHAPTER 28

Configuring FICON

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing in-band management of the switch from FICON processors.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations (see Chapter 39, "Configuring Fabric Binding"). The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an LIR to a registered Nx port.

This chapter includes the following sections:

About FICON

The FICON feature is not supported on:

- Cisco MDS 9120 switches
- Cisco MDS 9124 switches
- Cisco MDS 9140 switches
- The 32-port Fibre Channel switching module
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeSystem

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high availability platform. This solution simplifies purchasing, reduces deployment and management costs, and reduces the complex evolution to shared mainframe and open systems storage networks (see Figure 28-1).

Figure 28-1 Shared System Storage Network

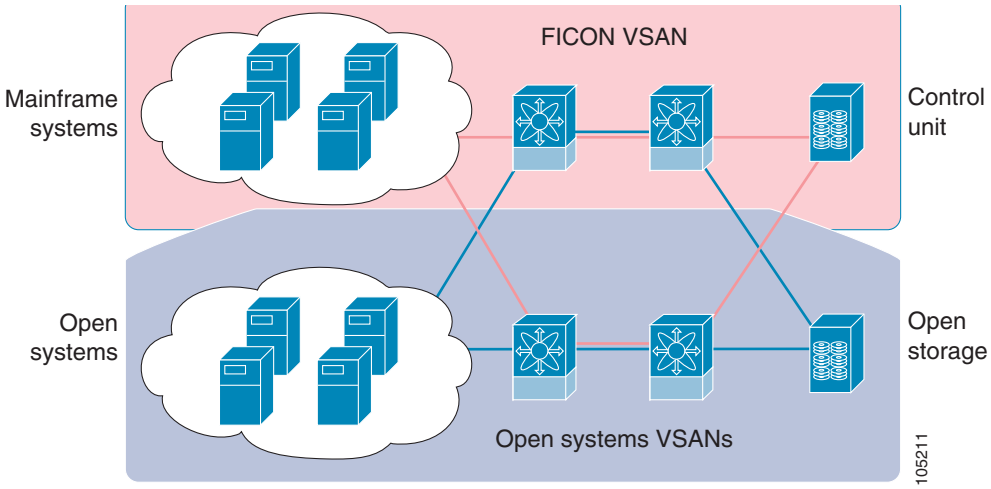

FCP and FICON are different FC4 protocols and their traffic is independent of each other. Devices using these protocols should be isolated using VSANs.

This section includes the following topics:

FICON Requirements

The FICON feature has the following requirements:

- You can implement FICON features in the following switches:
  - Any switch in the Cisco MDS 9500 Series.
  - Any switch in the Cisco MDS 9200 Series (including the Cisco MDS 9222i Multiservice Modular Switch).
  - Cisco MDS 9134 Multilayer Fabric Switch.
  - MDS 9000 Family 18/4-Port Multiservice Module.
- You need the MAINFRAME_PKG license to configure FICON parameters. To extendyour FICON configuration over a WAN link using FCIP, you need the appropriate SAN_EXTN_OVER_IP license for the module you are using. For more information, see Chapter 3, “Obtaining and Installing Licenses”.

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches and includes the following topics:

Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. Further, the ports in each island may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can introduce greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed. VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see Figure 28-2).

Figure 28-2 VSAN-Specific Fabric Optimization

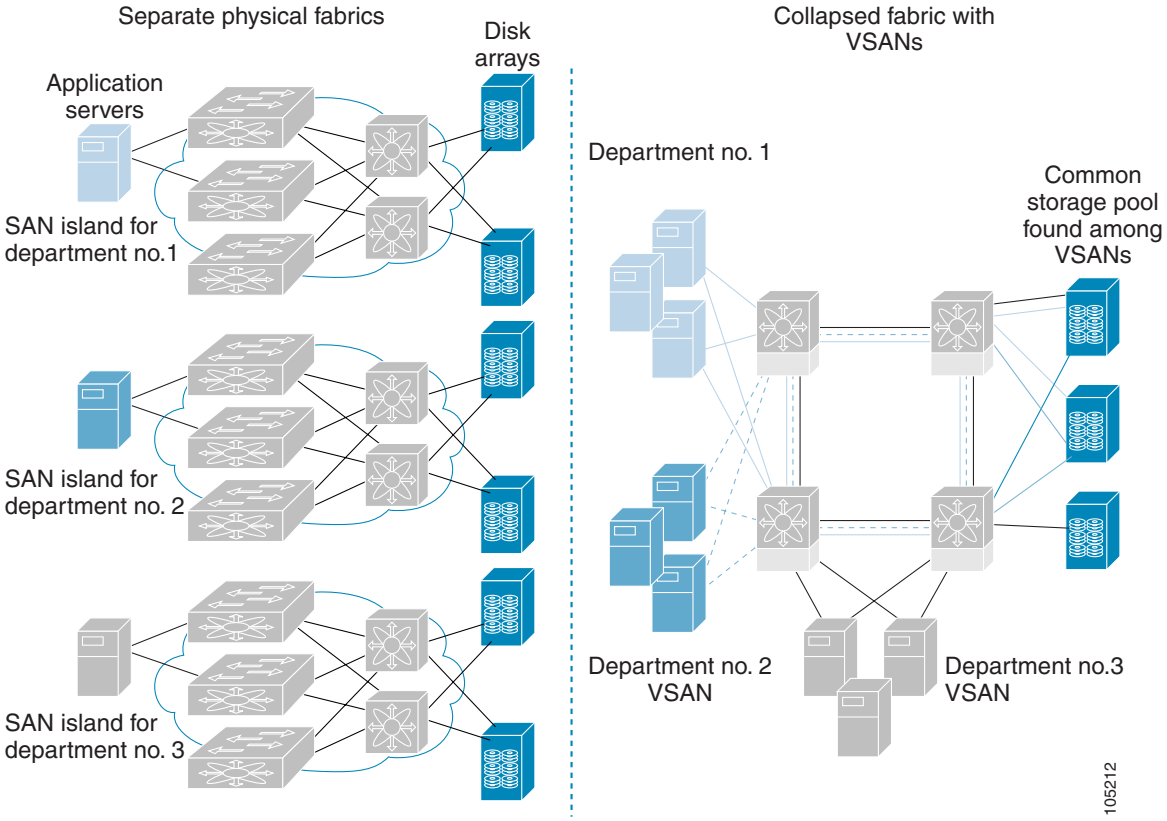

VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.

Note While you can configure VSANs in any Cisco MDS switch, you can enable FICON in at most eight of these VSANs. The number of VSANs configured depends on the platform.

Note Mainframe users can think of VSANs as being like FICON LPARs in the MDS SAN fabric. You can partition switch resources into FICON LPARs (VSANs) that are isolated from each other, in much the same way that you can partition resources on a zSeries or DS8000. Each VSAN has its own set of fabric services (such as fabric server and name server), FICON Control Unit Port, domain ID, Fabric Shortest Path First (FSPF) routing, operating mode, IP address, and security profile.

FICON LPARs can span line cards and are dynamic in size. For example, one FICON LPAR with 10 ports can span 10 different line cards. FICON LPARs can also include ports on more than one switch in a cascaded configuration. The consistent fairness of the Cisco MDS 9000 switching architecture means that “all ports are created equal”, simplifying provisioning by eliminating the “local switching” issues seen on other vendors’ platforms.

Addition of ports to a FICON LPAR is a non-disruptive process. The maximum number of ports for a FICON LPAR is 255 due to FICON addressing limitations.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and thus simplifies business continuance strategies.

See Chapter 40, “Configuring FCIP.”

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of Inter- switch Links (ISLs) necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

See Chapter 16, “Configuring PortChannels” for more information on PortChannels.

VSANs for FICON and FCP Mixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex mixed environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems Fibre Channel Protocol (FCP) fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based mixed schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Mixed environments are addressed by the Cisco SAN-OS software. The challenge of mixing FCP and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol mixing at the port level. If these protocols are mixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

Tip When creating a mixed environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

  Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide.*

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 528 autosensing, 4/2/1-Gbps, 10-Gbps, FICON or FCP ports in any combination in a single chassis. See Chapter 9, “Configuring High Availability.”

- Infrastructure protection—Common software releases provide infrastructure protection across all Cisco MDS 9000 platforms. See Chapter 7, “Software Images.”

- VSAN technology—The Cisco MDS 9000 Family provides VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON mixed support. See Chapter 19, “Configuring and Managing VSANs.”

- Port-level configurations—There are BB_credits, beacon mode, and port security for each port. See the “About Buffer-to-Buffer Credits” section on page 12-33, “Identifying the Beacon LEDs” section on page 12-17, and Chapter 15, “Configuring Trunking.”

- Alias name configuration—Provides user-friendly aliases instead of the WWN for switches and attached node devices. See Chapter 23, “Configuring and Managing Zones.”

- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS and TACACS+ authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN

zoning, read-only zones, and VSAN-based access control. See Chapter 33, "Configuring RADIUS and TACACS+" Chapter 37, "Configuring FC-SP and DHCHAP," and Chapter 39, "Configuring Fabric Binding."

- Traffic encryption—IPSec is supported over FCIP. You can encrypt FICON and Fibre Channel traffic that is carried over FCIP. See Chapter 36, "Configuring IPsec Network Security."
- Local accounting log—View the local accounting log to locate FICON events. See the "MSCHAP Authentication" section on page 33-34 and "Local AAA Services" section on page 33-35.
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the "CUP In-Band Management" section on page 28-41.
- Port address-based configurations—Configure port name, blocked or unblocked state, and the prohibit connectivity attributes can be configured on the ports. See the "Configuring FICON Ports" section on page 28-24.
- You can display the following information:
  - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
  - Nodes attached to ports.
  - Port performance and statistics.
- Configuration files—Store and apply configuration files. See the "FICON Configuration Files" section on page 28-32.
- FICON and Open Systems Management Server features if installed. —See the "VSANs for FICON and FCP Mixing" section on page 28-5.
- Enhanced cascading support—See the "CUP In-Band Management" section on page 28-41.
- Date and time—Set the date and time on the switch. See the "Allowing the Host to Control the Timestamp" section on page 28-21.
- Configure SNMP trap recipients and community names—See the "Configuring SNMP Control of FICON Parameters" section on page 28-22.
- Call Home configurations—Configure the director name, location, description, and contact person. See Chapter 54, "Configuring Call Home."
- Configure preferred domain ID, FC ID persistence, and principal switch priority—See Chapter 17, "Configuring Domain Parameters."
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol decoding, and network analysis tools as well as integrated Call Home capability for added reliability, faster problem resolution, and reduced service costs. See Chapter 52, "Monitoring Network Traffic Using SPAN."
- Configure R_A_TOV, E_D_TOV—— See the "Fibre Channel Time Out Values" section on page 29-3.
- Director-level maintenance tasks—Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis. See Chapter 59, "Monitoring System Processes and Logs."
- Port-level incident alerts—Display and clear port-level incident alerts. See the "Clearing RLIR Information" section on page 28-32.

FICON Cascading

The Cisco MDS SAN-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch (see Chapter 39, "Configuring Fabric Binding").

FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the "About the Default Zone" section on page 23-9.
- Enable in-order delivery on the VSAN. See Chapter 25, "Configuring Fibre Channel Routing Services and Protocols."
- Enable (and if required, configure) fabric binding on the VSAN. See Chapter 39, "Configuring Fabric Binding."
- Verify that conflicting persistent FC IDs do not exist in the switch. See Chapter 17, "Configuring Domain Parameters."
- Verify that the configured domain ID and requested domain ID match. See Chapter 17, "Configuring Domain Parameters."
- Add the CUP (area FE) to the zone, if you are using zoning. See the "CUP In-Band Management" section on page 28-41.

If any of these requirements are not met, the FICON feature cannot be enabled.

FICON Port Numbering

With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the *port number*. A maximum of 255 port numbers are available. You can use the following port numbering schemes:

- Default port numbers based on the chassis type
- Reserved port numbers

This section includes the following topics:

- FC ID Allocation, page 28-14

Note You must enable FICON on the switch before reserving FICON port number (see the About Enabling FICON on a VSAN, page 28-15).

Default FICON Port Numbering Scheme

Default FICON port numbers are assigned by the Cisco MDS SAN-OS software based on the module and the slot in the chassis. The first port in a switch always starts with a zero (0) (see Figure 28-3).

Figure 28-3 Default FICON Port Number in Numbering on the Cisco MDS 9000 Family9509 Switch

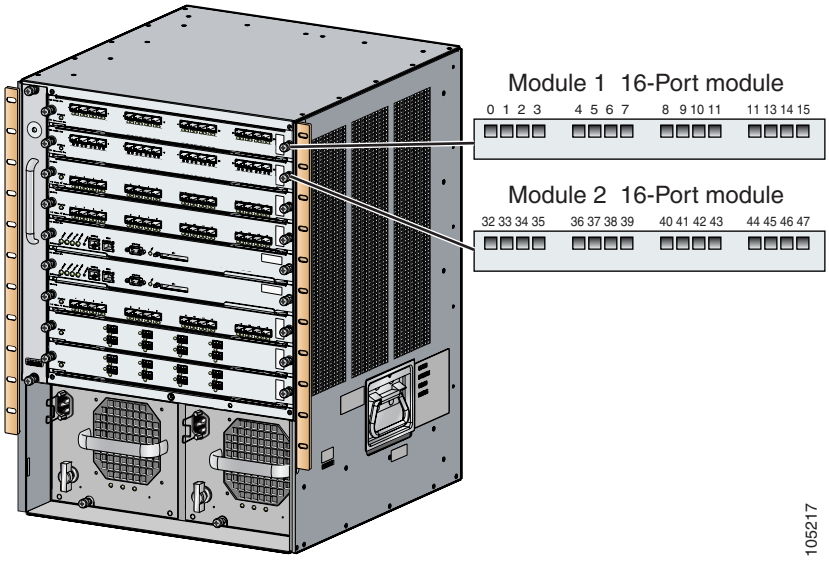

The default FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Thirty-two (32) port numbers are assigned to each slot on all Cisco MDS 9000 Family switches except for the Cisco MDS 9513 Director, which has 16 port numbers assigned for each slot. These default numbers are assigned regardless of the module's physical presence in the chassis, the port status (up or down), or the number of ports on the module (4, 12, 16, 24, or 48). If a module has fewer ports than the number of port numbers assigned to the slot, then the excess port numbers are unused. If a module has more ports than the number of port numbers assigned to the slot, the excess ports cannot be used for FICON traffic unless you manually assign port numbers.

Note You can use the **ficon slot assign port-numbers** command to make use of any excess ports by assigning numbers to the slots. Before doing this, however, we recommend that you review the default port number assignments for Cisco MDS 9000 switches shown inTable 28-1, and that you read the following sections to gain a complete understanding of FICON port numbering: "About the Reserved FICON Port Numbering Scheme" section on page 28-10, "FICON Port Numbering Guidelines" section on page 28-11, and "Assigning FICON Port Numbers to Slots" section on page 28-11.

Note Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

Table 28-3 lists the default port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 28-1 Default FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9200 Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	Similar to a switching module.
	Slot 2	32 through 63			
Cisco MDS 9222i Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated numbers.
	Slot 2	32 through 63			
Cisco MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	Supervisor modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			
Cisco MDS 9134 Director	Slot 1	0 through 33	34 through 59	60 through 253 and port 255	
Cisco MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			Supervisor modules are not allocated port numbers.
	Slot 6	None			
	Slot 7	128 through 159			The first 4, 12, 16, or 24 port numbers are used for a 4-port, 12-port,16-port, or 24-port module and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 8	160 through 191			
	Slot 9	192 through 223			

Table 28-1 Default FICON Port Numbering in the Cisco MDS 9000 Family (continued)

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9513 Director	Slot 1	0 through 15	224 through 249	250 through 253 and port 255	The first 4, 12 or 16 port numbers are used for a 4-port, 12-port or 16-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 2	16 through 31			
	Slot 3	32 through 47			
	Slot 4	48 through 63			
	Slot 5	64 through 79			
	Slot 6	80 through 95			
	Slot 7	None			Supervisor modules are not allocated port numbers.
	Slot 8	None			
	Slot 9	96 through 111			The first 4 or 12 port numbers are used for a 4-port or 12-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 10	112 through 127			
	Slot 11	128 through 143			
	Slot 12	144 through 159			
	Slot 13	160 through 175			

Port Addresses

By default, port numbers are the same as port addresses. You can swap the port addresses (see the "Port Swapping" section on page 28-36).

You can swap the port addresses by issuing the **ficon swap portnumber** command.

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is assigned by default to a slot in the chassis (see Table 28-3). An unimplemented port refers to any port address that is not assigned by default to a slot in the chassis (see Table 28-3).

About the Reserved FICON Port Numbering Scheme

A range of 250 port numbers are available for you to assign to all the ports on a switch. Table 28-3 shows that you can have more than 250 physical ports on a switch and the excess ports do not have port numbers in the default numbering scheme. When you have more than 250 physical ports on your switch, you can have ports without a port number assigned if they are not in a FICON VSAN, or you can assign duplicate port numbers if they are not used in the same FICON VSAN. For example, you can configure port number 1 on interface fc1/1 in FICON VSAN 10 and fc10/1 in FICON VSAN 20.

Note A VSAN can have a maximum of 250 port numbers.

Note FICON port numbers are not changed for ports that are active. You must first disable the interfaces using the **shutdown** command.

Note You can configure port numbers even when no module is installed in the slot.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—For example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—For example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- For slot 1, ports 0 to 31, or 0 to 15 have been assigned. Only the physical port fc1/5 with port number 4 is in VSAN 2. The rest of the physical ports are not in VSAN 2. The port numbers 0 to 249 are considered implemented for any FICON-enabled VSAN. Therefore, VSAN 2 has port numbers 0 to 249 and one physical port, fc1/4. The corresponding physical ports 0 to 3,and 5 to 249 are not in VSAN 2. When the FICON VSAN port address is displayed, those port numbers with the physical ports not in VSAN 2 are not installed (for example, ports 0 to 3, or 5 to 249).

  Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.
- The port is part of a PortChannel—For example, if interface fc 1/1 is part of PortChanne1 5, port address 0 is uninstalled in all FICON VSANs. See Table 28-3.

FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers do not change based on TE ports. Since TE ports appear in multiple VSANs, chassis-wide unique port numbers should be reserved for TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, then the associated ports will not come up.

  See the "About Port Numbers for FCIP and PortChannel" section on page 28-13.

Assigning FICON Port Numbers to Slots

You can use the **show ficon port-number assign** and **show ficon first-available port-number** commands to determine which port numbers to use.

Caution When you assign, change, or release a port number, the port reloads.

To assign FICON port numbers to a slot, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon slot 3 assign port-numbers 0-15, 48-63`	Reserves FICON port numbers 0 through 15 and 48 through 63 for up to 32 interfaces in slot 3.
	`switch(config)# ficon slot 3 assign port-numbers 0-15, 0-15`	Reserves FICON port numbers 0 through 15 for the first 16 interfaces and 0 through 15 for the next 16 interfaces in slot 3.
	`switch(config)# ficon slot 3 assign port-numbers 0-63`	Reserves FICON port numbers 0 through 63 for up to 64 interfaces in slot 3.
	`switch(config)# ficon slot 3 assign port-numbers 0-15, 56-63`	Changes the reserved FICON port numbers for up to 24 interfaces in slot 3.
	`switch(config)# no ficon slot 3 assign port-numbers 0-15, 56-63`	Releases the FICON port numbers.

Displaying the FICON Port Number Assignments

Use the **show ficon port-numbers assign** command to display the port numbers assigned on the switch.

```
switch# show ficon port-numbers assign
ficon slot 1 assign port-numbers 0-31
ficon slot 2 assign port-numbers 32-63
ficon slot 3 assign port-numbers 64-95
ficon slot 4 assign port-numbers 96-127
ficon logical-port assign port-numbers 128-153
```

Use the **show ficon port-numbers assign slot** command to display the port numbers assigned to a specific slot.

```
switch# show ficon port-numbers assign slot 2
ficon slot 2 assign port-numbers 32-63
```

Use the **show ficon port-numbers assign** command to display the port numbers reserved for logical ports.

```
switch# show ficon port-numbers assign logical-port
ficon logical-port assign port-numbers 128-153
```

About Port Numbers for FCIP and PortChannel

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

See the "Configuring FICON Ports" section on page 28-24 and the "Binding Port Numbers to FCIP Interfaces" section on page 28-25.

You can use the default port numbers if they are available (see Table 28-1 on page 28-9) or if you reserve port numbers from the pool of port numbers that are not reserved for Fibre Channel interfaces (see the"About the Reserved FICON Port Numbering Scheme" section on page 28-10).

To find the first available port number to bind an FCIP or PortChannel interface, use the **show ficon first-available port-number** command (see Example 28-12 on page 28-44).

Tip The **show ficon vsan portaddress brief** command displays the port number to interface mapping. You can assign port numbers in the PortChannel/FCIP range that are not already assigned to a PortChannel or FCIP interface (see Example 28-13 on page 28-44).

Reserving FICON Port Numbers for FCIP and PortChannel Interfaces

You must reserve port numbers for logical interfaces, such as FCIP and PortChannels, if you plan to use them.

To reserve FICON port numbers for logical interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon logical-port assign port-numbers 230-249`	Reserves port numbers 230 through 249 for FCIP and PortChannel interfaces.
Step 3	`switch(config)# ficon logical-port assign port-numbers 0xe6-0xf9`	Reserves port numbers 0xe6 through 0xf9 for FCIP and PortChannel interfaces. **Note** You cannot change port numbers that are active. You must disable the interfaces using the **shutdown** command and unbind port numbers using the **no ficon portnumber** command. See the "Configuring FICON Ports" section on page 28-24.
Step 4	`switch(config)# no ficon logical-port assign port-numbers 230-249`	Releases the port numbers. **Note** You cannot release port numbers for interfaces that are active.You must disable the interfaces using the **shutdown** command and unbind port numbers using the **no ficon portnumber** command. See the "Configuring FICON Ports" section on page 28-24.

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 .

Note You cannot configure persistent FC IDs in FICON-enabled VSANs.

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are shut down and restarted to switch from the dynamic to static FC IDs and vice versa (see Figure 28-4).

Figure 28-4 Static FC ID Allocation for FICON

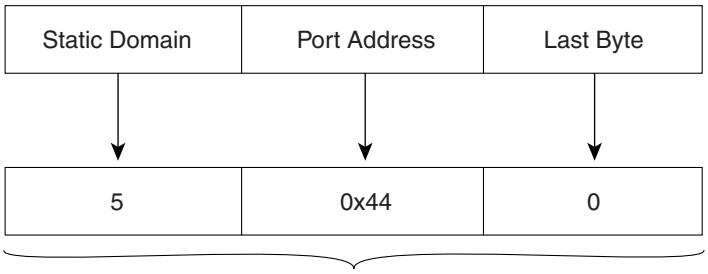

Static FC ID allocation for interface fc3/5 includes the static domain ID (5), the port address (0x44), and the last byte value (0).

Configuring FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on a per VSAN basis by using the Device Manager.

This section includes the following topics:

About Enabling FICON on a VSAN

By default FICON is disabled in all VSANs on the switch.

You can enable FICON on a per VSAN basis in one of the following ways:

- Use the automated **setup ficon** command.

  See the "Setting Up a Basic FICON Configuration" section on page 28-15.
- Manually addressing each prerequisite.

  See the "About FICON" section on page 28-1.
- Use Device Manager (refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*).

When you enable the FICON feature in Cisco MDS switches, the following apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.

  See the "About FICON Configuration Files" section on page 28-33.

Enabling and Disabling FICON on the Switch

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on the switch either explicitly or implicitly by enabling FICON on a VSAN. However, disabling FICON on all VSANs does not disable FICON on the switch. You must explicitly disable FICON.

To explicitly enable or disable FICON globally on the switch, following these steps.

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon enable`	Enables FICON globally on the switch.
Step 3	`switch(config)# no ficon enable`	Disables FICON globally on the switch and removes all FICON configuration.

Setting Up a Basic FICON Configuration

This section steps you through the procedure to set up FICON on a specified VSAN in a Cisco MDS 9000 Family switch.

Note Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.

Tip If you do not wish to answer a previously configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

To enable and set up FICON, follow these steps:

Step 1 Issue the **setup ficon** command at the EXEC command mode.

```
switch# setup ficon
                --- Ficon Configuration Dialog ---

This setup utility will guide you through basic Ficon Configuration
on the system.

Press Enter if you want to skip any dialog. Use ctrl-c at anytime
to skip all remaining dialogs.
```

Step 2 Enter **yes** (the default is **yes**) to enter the basic FICON configuration setup.

```
Would you like to enter the basic configuration dialog (yes/no) [yes]: yes
```

The FICON setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 3 Enter the VSAN number for which FICON should be enabled.

```
Enter vsan [1-4093]:2
```

Step 4 Enter **yes** (the default is **yes**) to create a VSAN.

```
vsan 2 does not exist, create it? (yes/no) [yes]: yes
```

Step 5 Enter **yes** (the default is **yes**) to confirm your VSAN choice:

```
Enable ficon on this vsan? (yes/no) [yes]: yes
```

Note At this point, the software creates the VSAN if it does not already exist.

Step 6 Enter the domain ID number for the specified FICON VSAN.

```
Configure domain-id for this ficon vsan (1-239):2
```

Step 7 Enter **yes** (the default is **no**) to set up FICON in cascaded mode. If you enter **no**, skip to Step 8 (see the "CUP In-Band Management" section on page 28-41).

```
Would you like to configure ficon in cascaded mode: (yes/no) [no]: yes
```

a. Assign the peer WWN for the FICON: CUP.

```
Configure peer wwn (hh:hh:hh:hh:hh:hh:hh:hh): 11:00:02:01:aa:bb:cc:00
```

b. Assign the peer domain ID for the FICON: CUP

```
Configure peer domain (1-239) :4
```

c. Enter **yes** if you wish to configure additional peers (and repeat Steps 7a and 7b). Enter **no**, if you do wish to configure additional peers.

```
Would you like to configure additional peers: (yes/no) [no]: no
```

Step 8 Enter **yes** (the default is **yes**) to allow SNMP permission to modify existing port connectivity parameters (see the "Configuring SNMP Control of FICON Parameters" section on page 28-22).

```
Enable SNMP to modify port connectivity parameters? (yes/no) [yes]: yes
```

Step 9 Enter **no** (the default is **no**) to allow the host (mainframe) to modify the port connectivity parameters, if required (see the "Allowing the Host to Change FICON Port Parameters" section on page 28-20).

```
Disable Host from modifying port connectivity parameters? (yes/no) [no]: no
```

Step 10 Enter **yes** (the default is **yes**) to enable the **active equals saved** feature (see the "Automatically Saving the Running Configuration" section on page 28-22).

```
Enable active=saved? (yes/no) [yes]: yes
```

Step 11 Enter **yes** (the default is **yes**) if you wish to configure additional FICON VSANs.

```
Would you like to configure additional ficon vsans (yes/no) [yes]: yes
```

Step 12 Review and edit the configuration that you have just entered.

Step 13 Enter **no** (the default is **no**) if you are satisfied with the configuration.

Note For documentation purposes, the following configurations shows three VSANs with different FICON settings. These settings provide a sample output for different FICON scenarios.

```
The following configuration will be applied:

fcdomain domain 2 static vsan 1
fcdomain restart disruptive vsan 1
fabric-binding database vsan 1
swwn 11:00:02:01:aa:bb:cc:00 domain 4
fabric-binding activate vsan 1
zone default-zone permit vsan 1
ficon vsan 1
no host port control

fcdomain domain 3 static vsan 2
fcdomain restart disruptive vsan 2
fabric-binding activate vsan 2 force
zone default-zone permit vsan 2
ficon vsan 2
no host port control
no active equals saved

vsan database
vsan 3
fcdomain domain 5 static vsan 3
fcdomain restart disruptive vsan 3
fabric-binding activate vsan 3 force
zone default-zone permit vsan 3
ficon vsan 3
no snmp port control
no active equals saved

Would you like to edit the configuration? (yes/no) [no]: no
```

Step 14 Enter **yes** (the default is **yes**) to use and save this configuration. The implemented commands are displayed. After FICON is enabled for the specified VSAN, you are returned to the EXEC mode switch prompt.

```
Use this configuration and apply it? (yes/no) [yes]: yes

`fcdomain domain 2 static vsan 1`
`fcdomain restart disruptive vsan 1`
`fabric-binding database vsan 1`
`swwn 11:00:02:01:aa:bb:cc:00 domain 4`
`fabric-binding activate vsan 1`
`zone default-zone permit vsan 1`
`ficon vsan 1`
`no host port control`

`fcdomain domain 3 static vsan 2`
`fcdomain restart disruptive vsan 2`
`fabric-binding activate vsan 2 force`
`zone default-zone permit vsan 2`
`ficon vsan 2`
`no host port control`
`no active equals saved`
```

Note If a new VSAN is created, two additional commands are displayed— **vsan database** and **vsan** *number*.

```
`vsan database`
`vsan 3`
`in-order-guarantee vsan 3`
`fcdomain domain 2 static vsan 3`
`fcdomain restart disruptive vsan 3`
`fabric-binding activate vsan 3 force`
`zone default-zone permit vsan 3`
`ficon vsan 3`
`no snmp port control`
Performing fast copy config...done.
switch#
```

Manually Enabling FICON on a VSAN

Note This section describes the procedure to manually enable FICON on a VSAN. If you have already enabled FICON on therequired VSAN using the automated setup (recommended), skip to the "Automatically Saving the Running Configuration" section on page 28-22.

To manually enable FICON on a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# vsan database`
	`switch(config-vsan-db)# vsan 5`
	`switch(config-vsan-db)# do show vsan usage`
	`4 vsan configured`
	`configured vsans:1-2,5,26`
	`vsans available for configuration:3-4,6-25,27-4093`
	`switch(config-vsan-db)# exit`
		Enables VSAN 5.

Step 3	`switch(config)# in-order-guarantee vsan 5`	Activates in-order delivery for VSAN 5. See Chapter 25, "Configuring Fibre Channel Routing Services and Protocols."
Step 4	`switch(config)# fcdomain domain 2 static vsan 2`	Configures the domain ID for VSAN 2. See Chapter 17, "Configuring Domain Parameters."
Step 5	`switch(config)# fabric-binding activate vsan 2 force`	Activates fabric binding on VSAN 2. See Chapter 39, "Configuring Fabric Binding."
Step 6	`switch(config)# zone default-zone permit vsan 2`	Sets the default zone to permit for VSAN 2. See the "CUP In-Band Management" section on page 28-41.
Step 7	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
	`switch(config)# no ficon vsan 6`	Disables the FICON feature on VSAN 6.
Step 8	`switch(config-ficon)# no host port control`	Prohibits mainframe users from moving the switch to an offline state. See the "Allowing the Host to Move the Switch Offline" section on page 28-20.

Configuring the code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.

Tip This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

To configure the **code-page** option in a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# code-page italy`	Configures the **italy** EBCDIC format.
	`switch(config-ficon)# no code-page`	Reverts to the factory default of using the **us-canada** EBCDIC format.

Allowing the Host to Move the Switch Offline

By default, hosts are allowed to move the switch to an offline state. To do this, the host sends "Set offline" command (x'FD') to CUP (Control Unit Port).

To allow the host to move the switch to an offline state, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# no host control switch offline`	Prohibits mainframe users from moving the switch to an offline state.
	`switch(config-ficon)# host control switch offline`	Allows the host to move the switch to an offline state (default) and shuts down the ports.

Allowing the Host to Change FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

Use the **host port control** command to permit mainframe users to configure FICON parameters.

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# no host port control`	Prohibits mainframe users from configuring FICON parameters on the Cisco MDS switch.
	`switch(config-ficon)# host port control`	Allows mainframe users to configure FICON parameters on the Cisco MDS switch (default).

Allowing the Host to Control the Timestamp

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS 9000 Family switch represents a virtual director. The clock and time present in each virtual director can be different.To maintain separate clocks for each VSAN, the Cisco SAN-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe) sets the time, the Cisco SAN-OS software updates this difference between the clocks. When a host reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

The VSAN-clock's current time is reported in the output of **show ficon vsan** *vsan-id*, **show ficon**, and **show accounting log** commands.

To configure host control of the timestamp, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# no host set-timestamp`	Prohibits mainframe users from changing the VSAN-specific clock.
	`switch(config-ficon)# host set-timestamp`	Allows the host to set the clock on this switch (default).

Clearing the Time Stamp

Note You can clear time stamps only from the Cisco MDS switch—not the mainframe.

Use the **clear ficon vsan** *vsan-id* **timestamp** command in EXEC mode to clear the VSAN clock.

```
switch# clear ficon vsan 20 timestamp
```

Configuring SNMP Control of FICON Parameters

By default, SNMP users can configure FICON parameters through the Cisco MDS 9000 Family Fabric Manager.

Note If you disable SNMP in the Cisco MDS switch, you cannot configure FICON parameters using the Fabric Manager.

To configure SNMP control of FICON parameters, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# no snmp port control`	Prohibits SNMP users from configuring FICON parameters.
	`switch(config-ficon)# snmp port control`	Allows SNMP users to configure FICON parameters (default).

About FICON Device Allegiance

FICON requires serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.

Caution This task discards the currently executing session.

Clearing FICON Device Allegiance

You can clear the current device allegiance by issuing the **clear ficon vsan** *vsan-id* **allegiance** command in EXEC mode.

```
switch# clear ficon vsan 1 allegiance
```

Automatically Saving the Running Configuration

Cisco MDS SAN-OS provides an option to automatically save any configuration changes to the startup configuration. This ensures that the new configuration is present after a switch reboot. The **active equals saved** option can be enable on any FICON VSAN.

Table 28-2 displays the results of the **active equals saved** command and the implicit **copy running-config startup-config** command in various scenarios.

If the **active equals saved** is enabled in any FICON-enabled VSAN in the fabric, then the following apply (see Number 1 and 2 in Table 28-2):

- All configuration changes (FICON-specific or not) are automatically saved to persistent storage (implicit **copy running start**) and stored in the startup configuration.
- FICON-specific configuration changes are immediately saved to the IPL file (see the "FICON Configuration Files" section on page 28-32).

If the **active equals saved** is not enabled in any FICON-enabled VSAN in the fabric, then FICON-specific configuration changes are not saved in the IPL file and an implicit **copy running startup** is not issued—you must issue the **copy running start** command explicitly (see number 3 in Table 28-2).

Table 28-2 Saving the Active FICON and Switch Configuration

Number	FICON-enabled VSAN?	`active equals saved` Enabled?	Implicit[1] `copy running start` Issued?	Notes
1	Yes	Yes (in all FICON VSANs)	Implicit	FICON changes written to the IPL file. Non-FICON changes saved to startup configuration and persistent storage.
2	Yes	Yes (even in one FICON VSAN)	Implicit	FICON changes written to IPL file for only the VSAN that has **active equals saved** option enabled. Non-FICON changes saved to startup configuration and persistent storage.
3	Yes	Not in any FICON VSAN	Not implicit	FICON changes are not written to the IPL file. Non-FICON changes are saved in persistent storage—only if you explicitly issue the **copy running start** command.
4	No	Not applicable		

1. When the Cisco SAN-OS software implicitly issues a **`copy running-config startup-config`** command in the Cisco MDS switch, only a binary configuration is generated—an ASCII configuration is not generated (see Example 28-24 on page 28-49). If you wish to generate an additional ASCII configuration at this stage, you must explicitly issue the **`copy running-config startup-config`** command again.

Note If **active equals saved** is enabled, the Cisco SAN-OS software ensures that you do not have to perform the **copy running startup** command for the FICON configuration as well. If your switch or fabric consists of multiple FICON-enabled VSANs, and one of these VSANs have **active equals saved** enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

To automatically save the running configuration, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.

	Command	Purpose
Step 3	`switch(config-ficon)# active equals saved`	Enables the automatic save feature for all VSANs in the switch or fabric.
	`switch(config-ficon)# no active equals saved`	Disables automatic save for this VSAN.

Configuring FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family of switches.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

This section includes the following topics:

Binding Port Numbers to PortChannels

Caution All port number assignments to PortChannels or FCIP interfaces are lost (cannot be retrieved) when FICON is disabled on all VSANs.

You can bind (or associate) a PortChannel with a FICON port number to bring up that interface.

To bind a PortChannel with a FICON port number, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface Port-channel 1` `switch(config-if)#`	Enters the PortChannel interface configuration mode.
Step 3	`switch(config-if)# ficon portnumber 234`	Assigns the FICON port number to the selected PortChannel port.

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) an FCIP interface with a FICON port number to bring up that interface.

To bind an FCIP interface with a FICON port number, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch(config-if)# ficon portnumber 208`	Assigns the FICON port number to the selected FCIP interface.

Configuring Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an Off-line state (OLS) primitive sequence on a blocked port.

Caution You cannot block or prohibit the CUP port (0XFE).

If a port is shut down, unblocking that port does not initialize the port.

Note The **shutdown/no shutdown** port state is independent of the **block/no block** port state.

To block or unblock port addresses in a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# portaddress 1 - 5` `switch(config-ficon-portaddr)#`	Selects port address 1 to 5 for further configuration.
Step 4	`switch(config-ficon-portaddr)# block`	Disables a range of port addresses and retains it in the operationally down state.
	`switch(config-ficon-portaddr)# no block`	Enables the selected port address and reverts to the factory default of the port address not being blocked.

Port Prohibiting

To prevent implemented ports from talking to each other, configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.

Tip You cannot prohibit a PortChannel or FCIP interface.

Unimplemented ports are always prohibited. In addition, prohibit configurations are always symmetrically applied—if you prohibit port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.

Note If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode or in TE mode.

Configuring the Default State for Port Prohibiting

By default, port prohibiting is disabled on the implemented interfaces on the switch. As of Cisco MDS SAN-OS Release 3.0(2), you can change the default port prohibiting state to enabled in VSANs that you create and then selectively disable port prohibiting on implemented ports, if desired. Also, only the FICON configuration files created after you change the default have the new default setting (see the "FICON Configuration Files" section on page 28-32).

To change the default port prohibiting setting for all implemented interfaces on the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon port default-state prohibit-all`	Enables port prohibiting as the default for all implemented interfaces on the switch.
	`switch(config)# no ficon port default-state prohibit-all`	Disables (default) port prohibiting as the default for all implemented interfaces on the switch.

Use the **show ficon port default-state** command to display the port prohibiting default state configuration.

```
switch# show ficon port default-state
Port default state is prohibit-all
```

Configuring Port Prohibiting

To prohibit port addresses in a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# portaddress 7` `switch(config-ficon-portaddr)#`	Selects port address 7 for further configuration.

	Command	Purpose
Step 4	`switch(config-ficon-portaddr)# prohibit portaddress 3-5`	Prohibits port address 7 in VSAN 2 from talking to ports 3, 4, and 5.
	`switch(config-ficon-portaddr)# no prohibit portaddress 5`	Removes port address 5 from a previously prohibited state.

Assigning a Port Address Name

To assign a port address name, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# portaddress 7` `switch(config-ficon-portaddr)#`	Selects port address 7 for further configuration.
Step 4	`switch(config-ficon-portaddr)# name SampleName`	Assigns a name to the port address. **Note** The port address name is restricted to 24 alphanumeric characters.
	`switch(config-ficon-portaddr)# no name SampleName`	Deletes a previously configured port address name.

About RLIR

The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an Link Incident Record (LIR) to a registered Nx port. It is a highly available application.

When an LIR is detected in FICON-enabled switches in the Cisco MDS 9000 Family from a RLIR Extended Link Service (ELS), the switch sends that record to the members in its Established Registration List (ERL).

In case of multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter-Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends it to the members of the ERL.

The Nx ports interested in receiving the RLIR ELS send the Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

The RLIR data is written to persistent storage when the **copy running-config startup-config** command is issued.

Specifying an RLIR Preferred Host

As of Cisco MDS SAN-OS Release 3.0(3), you can specify a preferred host to receive RLIR frames. The MDS switch sends RLIR frames to the preferred host only if it meets the following conditions:

- No host in the VSAN is registered for RLIR with the registration function set to “always receive.” If one or more hosts in the VSAN are registered as “always receive,” then RLIR sends only to these hosts and not to the configured preferred host.

- The preferred host is registered with the registration function set to “conditionally receive.”

Note If all registered hosts have the registration function set to “conditionally receive,” then the preferred host receives the RLIR frames.

You can specify only one RLIR preferred host per VSAN. By default, the switch sends RLIR frames to one of the hosts in the VSAN with the register function set to “conditionally receive” if no hosts have the register function set to “always receive.”

To specify the RLIR preferred host for a VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rlir preferred-cond fcid 0x772c00 vsan 5`	Specifies FC ID 0x772c00 as the RLIR preferred host in VSAN 5. (FC ID 0x772c00 is used here as an example.)
	`switch(config)# no rlir preferred-cond fcid 0x654321 vsan 2`	Removes FC ID 0x772c00 as the RLIR preferred host for VSAN 5.

To display the RLIR preferred host configuration, use the **show rlir erl** command.

```
switch# show rlir erl
Established Registration List for VSAN: 5
-----------------------------------------------
FC-ID LIRR FORMAT REGISTERED FOR
-----------------------------------------------
0x772c00 0x18 conditional receive(*)
0x779600 0x18 conditional receive
0x779700 0x18 conditional receive
0x779800 0x18 conditional receive
Total number of entries = 4
(*) - Denotes the preferred host
```

Displaying RLIR Information

The show rlir statistics command displays the complete statistics of LIRR, RLIR, and DRLIR frames. It lists the number of frames received, sent, and rejected. Specify the VSAN ID to obtain VSAN statistics for a specific VSAN. If you do not specify the VSAN ID, then the statistics are shown for all active VSANs (see Examples 28-1 and 28-2).

Example 28-1 Displays RLIR Statistics for All VSANs

```
switch# show rlir statistics

Statistics for VSAN: 1
------------------------

Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
```

```
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0

Statistics for VSAN: 100
------------------------

Number of LIRR received      = 26
Number of LIRR ACC sent      = 26
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 815
Number of RLIR ACC received  = 815
Number of RLIR RJT received  = 0
Number of DRLIR received     = 417
Number of DRLIR ACC sent     = 417
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 914
Number of DRLIR ACC received = 828
Number of DRLIR RJT received = 0
```

Example 28-2 Displays RLIR Statistics for a Specified VSAN

```
switch# show rlir statistics vsan 4

Statistics for VSAN: 4
------------------------

Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0
```

The **show rlir erl** command shows the list of Nx ports that are registered to receive the RLIRs with the switch. If the VSAN ID is not specified, the details are shown for all active VSANs (see Examples 28-3 and 28-4).

Example 28-3 Displays All ERLs

```
switch# show rlir erl

Established Registration List for VSAN: 2
-------------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-------------------------------------------
0x0b0200    0x18           always receive
Total number of entries = 1

Established Registration List for VSAN: 100
-------------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-------------------------------------------
0x0b0500    0x18           conditional receive
0x0b0600    0x18           conditional receive
Total number of entries = 2
```

In Example 28-3, if the `Registered For` column states that an FC ID is `conditional receive`, the source port is registered as a valid recipient of subsequent RLIRs. This source port is selected as an RLIR recipient only if no other ERL recipient is selected.

In Example 28-3, if the `Registered For` column states that an FC ID is `always receive`, the source port is registered as a valid recipient of subsequent RLIRs. This source port is always selected as an LIR recipient.

Note If an *always receive* RLIR is not registered for any N port or if the delivery of an RLIR fails for one of those ports, then the RLIR is sent to a port registered to *conditional receive* RLIRs.

Example 28-4 Displays ERLs for the Specified VSAN

```
switch# show rlir erl vsan 100
Established Registration List for VSAN: 100
-------------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-------------------------------------------
0x0b0500    0x18           conditional receive
0x0b0600    0x18           conditional receive

Total number of entries = 2
```

Note In Example 28-5, through Example 28-7, if the host time stamp (marked by the *) is available, it is printed along with the switch time stamp. If the host time stamp is not available, only the switch time stamp is printed.

Example 28-5 Displays the LIR History

```
switch# show rlir history

Link incident history
-------------------------------------------------------------------------------
*Host Time Stamp
 Switch Time Stamp          Port   Interface   Link Incident
-------------------------------------------------------------------------------

*Sun Nov 30 21:47:28 2003
```

```
 Sun Nov 30 13:47:55 2003     2        fc1/2   Implicit Incident
*Sun Nov 30 22:00:47 2003
 Sun Nov 30 14:01:14 2003     2        fc1/2   NOS Received
*Sun Nov 30 22:00:55 2003
 Sun Nov 30 14:01:22 2003     2        fc1/2   Implicit Incident
*Mon Dec 1 20:14:26 2003
 Mon Dec  1 12:14:53 2003     4        fc1/4   Implicit Incident
*Mon Dec 1 20:14:26 2003
 Mon Dec  1 12:14:53 2003     4        fc1/4   Implicit Incident
*Thu Dec 4 04:43:32 2003
 Wed Dec  3 20:43:59 2003     2        fc1/2   NOS Received
*Thu Dec 4 04:43:41 2003
 Wed Dec  3 20:44:08 2003     2        fc1/2   Implicit Incident
*Thu Dec 4 04:46:53 2003
 Wed Dec  3 20:47:20 2003     2        fc1/2   NOS Received
*Thu Dec 4 04:47:05 2003
 Wed Dec  3 20:47:32 2003     2        fc1/2   Implicit Incident
*Thu Dec 4 04:48:07 2003
 Wed Dec  3 20:48:34 2003     2        fc1/2   NOS Received
*Thu Dec 4 04:48:39 2003
 Wed Dec  3 20:49:06 2003     2        fc1/2   Implicit Incident
*Thu Dec 4 05:02:20 2003
 Wed Dec  3 21:02:47 2003     2        fc1/2   NOS Received
...
```

Example 28-6 Displays Recent LIRs for a Specified Interface

```
switch# show rlir recent interface fc1/1-4

Recent link incident records
-------------------------------------------------------------------------------
Host Time Stamp          Switch Time Stamp        Port Intf   Link Incident
-------------------------------------------------------------------------------
Thu Dec 4 05:02:29 2003  Wed Dec 3 21:02:56 2003   2   fc1/2  Implicit Incident
Thu Dec 4 05:02:54 2003  Wed Dec 3 21:03:21 2003   4   fc1/4  Implicit Incident
```

Example 28-7 Displays Recent LIRs for a Specified Port Number

```
switch# show rlir recent portnumber 1-4

Recent link incident records
-------------------------------------------------------------------------------
Host Time Stamp          Switch Time Stamp        Port Intf    Link Incident
-------------------------------------------------------------------------------
Thu Dec 4 05:02:29 2003  Wed Dec 3 21:02:56 2003   2   fc1/2   Implicit Incident
Thu Dec 4 05:02:54 2003  Wed Dec 3 21:03:21 2003   4   fc1/4   Implicit Incident
```

As of Cisco SAN-OS Release 3.0(3), the **show rlir history** command output includes remote link incidents that are received as DRLIRs from other switches. RLIRs are generated as a result of DRLIRs as in previous Cisco SAN-OS releases (see Example 28-8).

Example 28-8 Displays the LIR History as of Cisco SAN-OS Release 3.0(3)

```
switch# show rlir history

Link incident history
-------------------------------------------------------------------------------------------------------------------
---
```

```
 Host Time Stamp         Switch Time Stamp       VSAN   Domain   Port    Intf        Link Incident
Loc/Rem
-------------------------------------------------------------------------------------------------------
---
 Sep 20 12:42:44 2006    Sep 20 12:42:44 2006    ****    ****    0x0b    fc1/12      Loss of sig/sync     LOC
 Reported Successfully to: [0x640001] [0x640201]
 Sep 20 12:42:48 2006    Sep 20 12:42:48 2006    ****    ****    0x0b    fc1/12      Loss of sig/sync     LOC
 Reported Successfully to: [0x640001] [0x640201]
 *** ** **:**:** ****    Sep 20 12:42:51 2006    1001    230     0x12    ****        Loss of sig/sync     REM
 Reported Successfully to: [0x640001] [0x640201]
 Sep 20 12:42:55 2006    Sep 20 12:42:55 2006    ****    ****    0x0b    fc1/12      Loss of sig/sync     LOC
 Reported Successfully to: None [No Registrations]
 *** ** **:**:** ****    Sep 20 12:45:56 2006    1001    230     0x12    ****        Loss of sig/sync     REM
 Reported Successfully to: None [No Registrations]
 *** ** **:**:** ****    Sep 20 12:45:56 2006    1001    230     0x12    ****        Loss of sig/sync     REM
 Reported Successfully to: None [No Registrations]
 Sep 20 12:52:45 2006    Sep 20 12:52:45 2006    ****    ****    0x0b    fc1/12      Loss of sig/sync     LOC
 Reported Successfully to: None [No Registrations]

**** - Info not required/unavailable
```

Clearing RLIR Information

Use the **clear rlir statistics** command to clear all existing statistics for a specified VSAN.

```
switch# clear rlir statistics vsan 1
```

Use the **clear rlir history** command to clear the RLIR history where all link incident records are logged for all interfaces.

```
switch# clear rlir history
```

Use the **clear rlir recent interface** command to clear the most recent RLIR information for a specified interface.

```
switch# clear rlir recent interface fc 1/2
```

Use the **clear rlir recent portnumber** command to clear the most recent RLIR information for a specified port number.

```
switch# clear rlir recent portnumber 16
```

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or Fabric Manager applications to operate on these FICON configuration files.

Note Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.

Caution When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name

Note Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, static domain ID configuration, and fabric binding configuration.

See the "Managing Configuration Files" section on page 8-1 for details on the normal configuration files used by Cisco MDS switches.

This section includes the following topics:

About FICON Configuration Files

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco SAN-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Applying the Saved Configuration Files to the Running Configuration

You can apply the configuration from the saved files to the running configuration using the **ficon vsan** *number* **apply file** *filename* command.

```
switch# ficon vsan 2 apply file SampleFile
```

Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.

To edit the contents of a specified FICON configuration file, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ficon vsan 2` `switch(config-ficon)#`	Enables FICON on VSAN 2.
Step 3	`switch(config-ficon)# file IplFile1` `switch(config-ficon-file)#`	Accesses the FICON configuration file called IplFile1 for VSAN 2. If this file does not exist, it is created. **Note** All FICON file names are restricted to eight alphanumeric characters.
	`switch(config-ficon)# no file IplFileA`	Deletes a previously created FICON configuration file.
Step 4	`switch(config-ficon-file)# portaddress 3` `switch(config-ficon-file-portaddr)#`	Enters the submode for port address 3 to edit the contents of the configuration file named IplFile1. **Note** The running configuration is not applied to the current configuration. The configuration is only applied when the **ficon vsan** *number* **apply file** *filename* command is issued.
Step 5	`switch(config-ficon-file-portaddr)# prohibit portaddress 5`	Edits the content of the configuration file named IplFile1 by prohibiting port address 5 from accessing port address 3.
Step 6	`switch(config-ficon-file-portaddr)# block`	Edits the content of the configuration file named IplFile1 by blocking a range of port addresses and retaining them in the operationally down state.
Step 7	`switch(config-ficon-file-portaddr)# name P3`	Edits the content of the configuration file named IplFile1 by assigning the name P3 to port address 3. If the name did not exist, it is created. If it existed, it is overwritten.

Displaying FICON Configuration Files

Use the **show ficon vsan** *vsan-id* **file all** command to display the contents of all FICON configuration files.

```
switch# show ficon vsan 2 file all
File IPL      is locked
FICON configuration file IPLFILEA in vsan 2
Description:
    Port address 0(0)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)

    Port address 1(0x1)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)
Port address 2(0x2)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)

    Port address 3(0x3)
        Port name is P3
        Port is blocked
        Prohibited port addresses are 5,250-253,255(0x5,0xfa-0xfd,0xff)
...
```

Use the **show ficon vsan** *vsan-id* **file name** command to display the contents of a specific FICON configuration file.

```
switch# show ficon vsan 2 file name IPLfilea
FICON configuration file IPLFILEA in vsan 2
Description:
    Port address 0(0)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)

    Port address 1(0x1)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)

    Port address 2(0x2)
        Port name is
        Port is not blocked
        Prohibited port addresses are 250-253,255(0xfa-0xfd,0xff)

    Port address 3(0x3)
        Port name is P3
        Port is blocked
        Prohibited port addresses are 5,250-253,255(0x5,0xfa-0xfd,0xff)
```

Use the **show ficon vsan** *vsan-id* **file name** *filename* **portaddress** command to display the FICON configuration file information for a specific FICON port.

```
switch# show ficon vsan 2 file name IPLfilea portaddress 3
FICON configuration file IPLFILEA in vsan 2
Description:
    Port address 3(0x3)
        Port name is P3
        Port is blocked
        Prohibited port addresses are 5,250-253,255(0x5,0xfa-0xfd,0xff)
```

Copying FICON Configuration Files

Use the **ficon vsan** *vsan-id* **copy file** *existing-file-name save-as-file-name* command in EXEC mode to copy an existing FICON configuration file.

```
switch# ficon vsan 20 copy file IPL IPL3
```

You can see the list of existing configuration files by issuing the **show ficon vsan** *vsan-id* command.

```
switch# show ficon vsan 20
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Disabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Enabled
  Number of implemented ports are 250
  Key Counter is 5
  FCID last byte is 0
  Date/Time is same as system time (Wed Dec 3 20:10:45.924591 2003)
  Device Allegiance not locked
  Codepage is us-canada
  Saved configuration files
    IPL
    IPL3
```

Port Swapping

The FICON port swapping feature is only provided for maintenance purposes.

The FICON port swapping feature causes all configuration associated with *old-port-number* and *new port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for nonexistent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.
- If you swap a port in a module that has unlimited oversubscription ratios enabled with a port in a module that has limited oversubscription ratios, then you may experience a degradation in bandwidth.

Tip If **active equals saved** is enabled on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.

If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.

The **ficon swap portnumber** command is only associated with the two ports concerned. You must issue this VSAN-independent command from EXEC mode. Cisco MDS SAN-OS checks for duplicate port numbers in a VSAN before performing the port swap.

If you attempt to bring the port up by specifying the **ficon swap portnumber** *old-port-number new-port-number* **after swap noshut** command, you must explicitly issue the **no shutdown** command to resume traffic.

This section includes the following topics:

About Port Swapping

Be sure to follow these guidelines when using the FICON port swapping feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check to verify the extended BB_credits configuration.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values.
- Port tracking information is not included in port swapping. This information must be configured separately (see Chapter 57, “Configuring Port Tracking”).

Note The 32-port module guidelines also apply for port swapping configurations (see the “Fibre Channel Interfaces” section on page 12-1).

Swapping Ports

If there are no duplicate port numbers on the switch, you can swap physical Fibre Channel ports, except the port numbers, by following these steps:

Step 1 Issue the **ficon swap portnumber** *old-port-number new-port-number* command in EXEC mode.

Note The **ficon swap portnumber** command might fail if more than one interface on the MDS switch has the same port number as the *old-port-number* or *new-port-number* specified in the command.

The specified ports are operationally shut down.

Step 2 Physically swap the front panel port cables between the two ports.

Step 3 Issue the **no shutdown** command on each port to enable traffic flow.

Note If you specify the **ficon swap portnumber** *old-port-number new-port-number* **after swap noshut** command, the ports are automatically initialized.

If there are duplicate port numbers on the switch, you can swap physical Fibre Channel ports, including the port numbers, by following these steps:

Step 1 Issue the **ficon swap interface** *old-interface new-interface* command in EXEC mode.

The specified interfaces are operationally shut down.

Step 2 Physically swap the front panel port cables between the two ports.

Step 3 Issue the **no shutdown** command on each port to enable traffic flow.

Note If you specify the **ficon swap interface** *old-interface new-interface* **after swap noshut** command, the ports are automatically initialized.

FICON Tape Acceleration

The sequential nature of tape devices causes each I/O operation to the tape device over an FCIP link to incur the latency of the FCIP link. Throughput drastically decreases as the round-trip time through the FCIP link increases, leading to longer backup windows. Also, after each I/O operation, the tape device is idle until the next I/O arrives. Starting and stopping of the tape head reduces the lifespan of the tape, except when I/O operations are directed to a virtual tape.

Cisco MDS SAN-OS software provides acceleration for the following FICON tape write operations:

- The link between mainframe and native tape drives (both IBM and Sun/STK)
- The back-end link between the VSM (Virtual Storage Management) and tape drive (Sun/STK)

FICON tape acceleration over FCIP provides the following advantages:

- Efficiently utilizes the tape device by decreasing idle time
- More sustained throughput as latency increases
- Similar to FCP tape acceleration, and does not conflict with it

Note FICON tape read acceleration over FCIP is not supported.

Figure 28-5 through Figure 28-8 show supported configurations:

Figure 28-5 Host Directly Accessing IBM/STK (StorageTek) Library

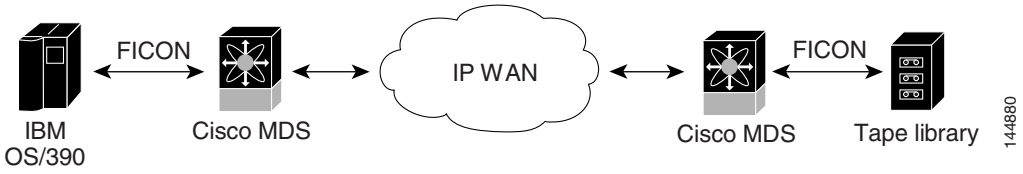

Figure 28-6 Host Accessing Standalone IBM-VTS (Virtual Tape Server) /STK-VSM (Virtual Shared Memory)

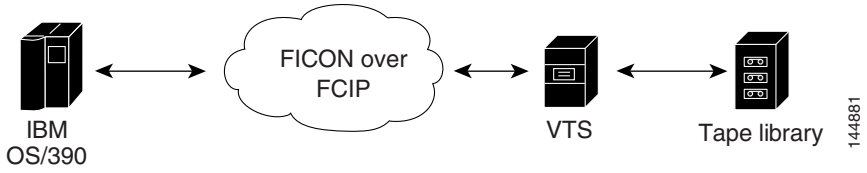

Figure 28-7 Host Accessing Peer-to-Peer VTS (Virtual Tape Server)

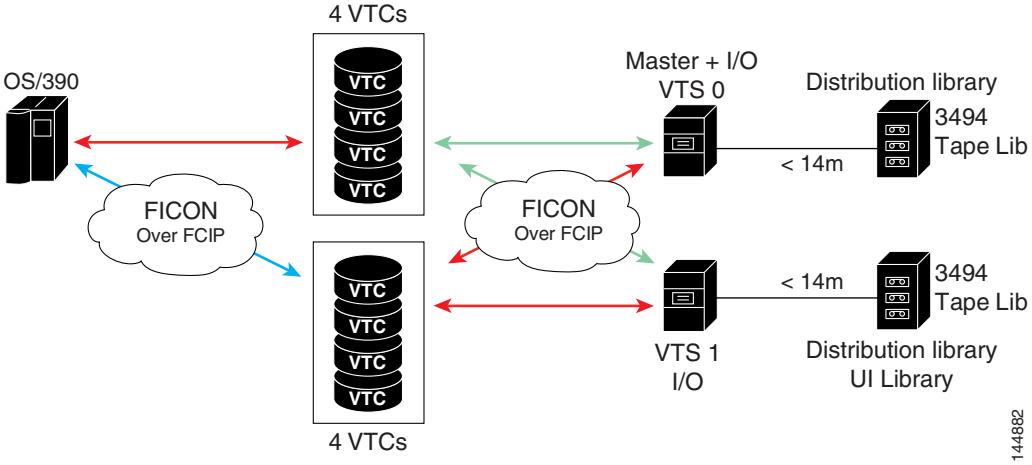

Figure 28-8 Host Accessing Peer-to-Peer VTS (Virtual Tape Server)

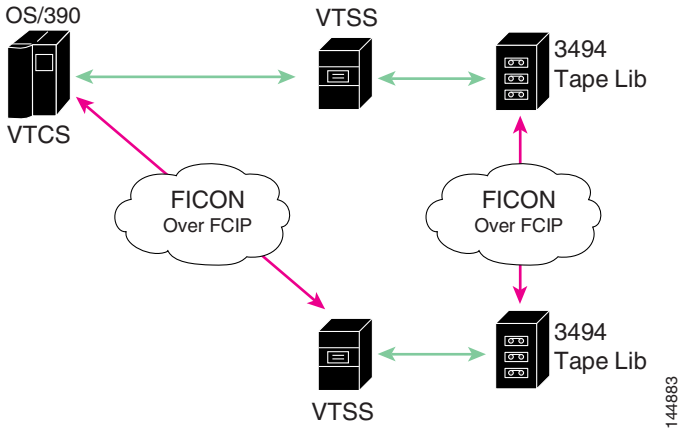

Note For information about FCIP tape acceleration, see "FCIP Tape Acceleration" section on page 40-29.

Configuring FICON Tape Acceleration

FICON tape acceleration has the following configuration considerations:

- In addition to the normal FICON configuration, FICON tape acceleration must be enabled on both ends of the FCIP interface. If only one end has FICON tape acceleration enabled, acceleration does not occur.
- FICON tape acceleration is enabled on a per VSAN basis.
- FICON tape acceleration cannot function if multiple ISLs are present in the same VSAN (PortChannels or FSPF load balanced).
- You can enable both Fibre Channel write acceleration and FICON tape acceleration on the same FCIP interface.
- Enabling or disabling FICON tape acceleration disrupts traffic on the FCIP interface.

To configure FICON tape acceleration, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fcip 2` `switch(config-if)#`	Specifies an FCIP interface and enters interface configuration submode.
Step 3	`switch(config-if)# ficon-tape-accelerator vsan 100` `This configuration change will disrupt all traffic on the FCIP interface in all VSANs. Do you wish to continue? [no] y`	Enables FICON tape acceleration over an FCIP interface.
	`switch(config-if)# no ficon-tape-accelerator vsan 100` `This configuration change will disrupt all traffic on the FCIP interface in all VSANs. Do you wish to continue? [no] y`	Disables (default) FICON tape acceleration over an FCIP interface.

Use the **show running-config** command to verify the FICON tape acceleration over FCIP configuration.

```
switch# show running-config | begin "interface fcip"
interface fcip2
  ficon-tape-accelerator vsan 100
  no shutdown
...
```

Moving a FICON VSAN to an Offline State

Issue the **ficon vsan** *vsan-id* **offline** command in EXEC mode to log out all ports in the VSAN that need to be suspended.

Issue the EXEC-level **ficon vsan** *vsan-id* **online** command in EXEC mode to remove the offline condition and to allow ports to log on again.

Note This command can be issued by the host if the host is allowed to do so (see the "Allowing the Host to Move the Switch Offline" section on page 28-20).

CUP In-Band Management

The Control Unit Port (CUP) protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.

Note The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

This section includes the following topics:

Placing CUPs in a Zone

To place the CUP in a zone, follow these steps:

Step 1 Set the default zone to permit for the required VSAN.

```
switch# config t
switch(config)# zone default-zone permit vsan 20
```

Step 2 Issue the **show fcns database** command for the required FICON CUP WWN.

```
switch# show fcns database vsan 20

VSAN 20:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x0d0d00    N     50:06:04:88:00:1d:60:83 (EMC)           FICON:CU
0x0dfe00    N     25:00:00:0c:ce:5c:5e:c2 (Cisco)         FICON:CUP
0x200400    N     50:05:07:63:00:c2:82:d3 (IBM)           scsi-fcp FICON:CU f..
0x200800    N     50:05:07:64:01:40:15:0f (IBM)           FICON:CH
0x20fe00    N     20:00:00:0c:30:ac:9e:82 (Cisco)         FICON:CUP

Total number of entries = 5
```

Note If more than one `FICON:CUP WWN` exists in this fabric, be sure to add all the `FICON:CUP WWN` PWWNs to the required zone. The previous sample output displays multiple `FICON:CUP` occurrences to indicate a cascade configuration.

Step 3 Add the identified FICON:CUP WWN to the zone database.

```
switch(config)# zone name Zone1 vsan 20
switch(config-zone)# member pwwn 25:00:00:0c:ce:5c:5e:c2
```

Displaying Control Unit Information

Example 28-9 displays configured control device information.

Example 28-9 Displays Control Unit Information

```
switch# show ficon control-device sb3
Control Unit Image:0x80b9c2c
VSAN:20 CU:0x20fe00 CUI:0 CUD:0 CURLP:(nil)
ASYNC LP:(nil) MODE:1 STATE:1 CQ LEN:0 MAX:0
PRIMARY LP:   VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
ALTERNATE LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
Logical Path:0x80b9fb4
VSAN:20 CH:0x200600 CHI:15 CU:0x20fe00 CUI:0 STATE:1 FLAGS:0x1
LINK: OH:0x0 OC:0x0 IH:0x0 IC:0x0
DEV:  OH:0x0 OC:0x0 IH:0x0 IC:0x0
SENSE: 00 00 00 00 00 00 00 46
       30 20 00 00 00 00 00 00
       00 00 00 00 00 00 00 00
```

```
        00 00 00 00 00 00 00 00
IUI:0x0 DHF:0x0 CCW:0x0 TOKEN:0x0 PCCW:0x0 FCCW:0x0 PTOKEN:0x0 FTOKEN:0x0
CMD:0x0 CCW_FLAGS:0x0 CCW_COUNT:0 CMD_FLAGS:0x0 PRIO:0x0 DATA_COUNT:0
STATUS:0x0 FLAGS:0x0 PARAM:0x0 QTP:0x0 DTP:0x0
CQ LEN:0 MAX:0 DESTATUS:0x0
```

Displaying FICON Information

This section includes the following topics:

Receiving FICON Alerts

In Example 28-10 the `user alert mode is enabled` output confirms that you will receive an alert to indicate any changes in the FICON configuration.

Example 28-10 Displays Configured FICON Information

```
switch# show ficon
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Enabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Disabled
  Number of implemented ports are 250
  Key Counter is 73723
  FCID last byte is 0
  Date/Time is set by host to Sun Jun 26 00:04:06.991999 1904
  Device allegiance is locked by Host
  Codepage is us-canada
  Saved configuration files
    IPL
    _TSIRN00
```

Displaying FICON Port Address Information

Examples 28-11 to 28-14 display FICON Port Address information.

Example 28-11 Displays Port Address Information

```
switch# show ficon vsan 2 portaddress
Port Address 1 is not installed in vsan 2
    Port number is 1, Interface is fc1/1
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255

Port Address 2 is not installed in vsan 2
    Port number is 2, Interface is fc1/2
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255
...
Port Address 249 is not installed in vsan 2
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255

Port Address 250 is not installed in vsan 2
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255
```

Example 28-12 Displays the Available Port Numbers

```
switch# show ficon first-available port-number
Port number 129(0x81) is available
```

In Example 28-13, the interface column is populated with the corresponding interface if the port number is installed. If the port number is uninstalled, this space remains blank and indicates an unbound port number. For example, 56 is an unbound port number in Example 28-13.

Example 28-13 Displays Port Address Information in a Brief Format

```
switch# show ficon vsan 2 portaddress 50-55 brief
--------------------------------------------------------------------------------
Port    Port    Interface          Admin     Status           Oper    FCID
Address Number                     Blocked                    Mode
--------------------------------------------------------------------------------
50      50      fc2/18              on       fcotAbsent       --      --
51      51      fc2/19              off      fcotAbsent       --      --
52      52      fc2/20              off      fcotAbsent       --      --
53      53      fc2/21              off      fcotAbsent       --      --
54      54      fc2/22              off      notConnected     --      --
55      55      fc2/23              off      up               FL      0xea0000
56      56                          off      up               FL      0xea0000
```

Example 28-14 displays the counters in FICON version format 1 (32-bit format)

Example 28-14 Displays Port Address Counter Information

```
switch# show ficon vsan 20 portaddress 8 counters
Port Address 8(0x8) is up in vsan 20
    Port number is 8(0x8), Interface is fc1/8
    Version presented 1, Counter size 32b
    242811 frames input, 9912794 words
      484 class-2 frames, 242302 class-3 frames
      0 link control frames, 0 multicast frames
      0 disparity errors inside frames
      0 disparity errors outside frames
      0 frames too big, 0 frames too small
      0 crc errors, 0 eof errors
      0 invalid ordered sets
      0 frames discarded c3
      0 address id errors
    116620 frames output, 10609188 words
      0 frame pacing time
    0 link failures
    0 loss of sync
    0 loss of signal
    0 primitive seq prot errors
    0 invalid transmission words
    1 lrr input, 0 ols input, 5 ols output
    0 error summary
```

Displaying FICON Configuration File Information

Examples 28-15 to 28-17 display FICON configuration file information.

Example 28-15 Displays the Contents of the Specified FICON Configuration File

```
switch# show ficon vsan 3 file IPL
FICON configuration file IPL      in vsan 3
    Port address 1
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255

    Port address 2
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255

    Port address 3
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255

    Port address 4
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255

...
Port address 80
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255
```

```
    Port address 254
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,81-253,255
```

Example 28-16 Displays All FICON Configuration Files

```
switch# show ficon vsan 2
Ficon information for VSAN 2
  Ficon is enabled
  VSAN is active
  Host control is Enabled
  Host offline control is Enabled
  Clock alert mode is Disabled
  User alert mode is Disabled
  SNMP control is Disabled
  Active=Saved is Disabled
  Number of implemented ports are 250
  Key Counter is 9
  FCID last byte is 0
  Date/Time is same as system time(Sun Dec 14 01:26:30.273402 1980)
  Device Allegiance not locked
  Codepage is us-canada
  Saved configuration files
    IPL
    IPLFILE1
```

Example 28-17 Displays the Specified Port Addresses for a FICON Configuration File

```
switch# show ficon vsan 2 file iplfile1 portaddress 1-7
FICON configuration file IPLFILE1 in vsan 2
    Port address 1
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,241-253,255

    Port address 2
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,241-253,255

    Port address 3
        Port name is P3
        Port is not blocked
        Prohibited port addresses are 0,241-253,255
...
    Port address 7
        Port name is
        Port is not blocked
        Prohibited port addresses are 0,241-253,255
```

Displaying the Configured FICON State

If FICON is enabled on a VSAN, you can display the port address information for that VSAN (see Example 28-18).

Example 28-18 Displays the Specified Port Address When FICON Is Enabled

```
switch# show ficon vsan 2 portaddress 55
Port Address 55 is not installed in vsan 2
    Port number is 55, Interface is fc2/23
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255
    Admin port mode is FL
    Port mode is FL, FCID is 0xea0000
```

Displaying a Port Administrative State

Examples 28-19 to 28-20 display the administrative state of a FICON port. If the port is blocked, the **show ficon vsan** *number* **portaddress** *number* command displays the blocked state of the port. If a specific port is prohibited, this command also displays the specifically prohibited port (3) along with the ports that are prohibited by default (0, 241 to 253, and 255). If a name is assigned, that name is also displayed.

Example 28-19 Displays an Administratively Unblocked Port

```
switch# show ficon vsan 2 portaddress 2
Port Address 2(0x2) is not installed in vsan 2
    Port number is 2(0x2), Interface is fc1/2
    Port name is
    Port is not admin blocked
    Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
    Admin port mode is auto
    Peer is Unknown
```

Example 28-20 Displays an Administratively Blocked Port

```
switch# show ficon vsan 2 portaddress 1
Port Address 2(0x2) is not installed in vsan 2
    Port number is 2(0x2), Interface is fc1/2
    Port name is SampleName
    Port is admin blocked
    Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
    Admin port mode is auto
    Peer is Unknown
```

Displaying Buffer Information

In Example 28-21, the `Key Counter` column displays the 32-bit value maintained by Cisco MDS switches. This value is incremented when any port changes state in that VSAN. The key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

The director history buffer provides a mechanism to determine the change in the port state from the previous time when a value was contained in the key counter.

Example 28-21 Displays the History Buffer for the Specified VSAN

```
switch# show ficon vsan 20 director-history
Director History Buffer for vsan 20
-------------------------------------------
Key Counter           Ports Address
                      Changed
-------------------------------------------
74556                 43
74557                 44
74558                 45
74559                 46
74560                 47
74561                 48
74562                 49
74563                 50
74564                 51
74565                 52
74566                 53
74567                 54
74568                 55
74569                 56
74570                 57
74571                 58
74572                 59
74573                 60
74574                 61
74575                 62
74576                 63
74577                 64
74578
74579
74580                 1-3,5,10,12,14-16,34-40,43-45,47-54,56-57,59-64
74581                 3,5
74582                 64
74583
74584                 1-3,10,12,14-16,34-40,43-45,47-54,56-57,59-64
74585                 1
74586                 2
74587                 3
```

Displaying FICON Information in the Running Configuration

Example 28-22 displays the FICON-related information in the running configuration.

Example 28-22 Displays the Running Configuration Information

```
switch# show running-config
Building Configuration ...
in-order-guarantee
vsan database
  vsan 11 name "FICON11" loadbalancing src-dst-id
  vsan 75 name "FICON75" loadbalancing src-dst-id

fcdomain domain 11 static vsan 11
fcdomain domain 119 static vsan 75

fcdroplatency network 100 vsan 11
fcdroplatency network 500 vsan 75
```

```
fabric-binding enable
fabric-binding database vsan 11
  swwn 20:00:00:0d:ec:01:20:c0 domain 10
fabric-binding database vsan 75
  swwn 20:00:00:0d:ec:00:d6:40 domain 117
fabric-binding activate vsan 11
fabric-binding activate vsan 75

ficon vsan 75

interface port-channel 1
  ficon portnumber 0x80
  switchport mode E

snmp-server user mblair network-admin auth md5 0x688fa3a2e51ba5538211606e59ac292
7 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server user wwilson network-admin auth md5 0x688fa3a2e51ba5538211606e59ac29
27 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server host 171.71.187.101 traps version 2c public udp-port 1163
snmp-server host 172.18.2.247 traps version 2c public udp-port 2162

vsan database
  vsan 75 interface fc1/1
...
interface mgmt0
  ip address 172.18.47.39 255.255.255.128
  switchport speed 100
  switchport duplex full

no system health

ficon vsan 75
  file IPL
```

Displaying FICON Information in the Startup Configuration

Example 28-23 displays the FICON-related information in the startup configuration.

Example 28-23 Displays the Startup Configuration

```
switch# show startup-config
...
ficon vsan 2
file IPL
```

Example 28-24 displays the switch response to an implicitly-issued copy running start command. In this case, only a binary configuration is saved until you explicitly issue the **copy running start** command again (see Table 28-2)

Example 28-24 Displays the Startup Configuration Status

```
switch# show startup-config
No ASCII config available since configuration was last saved internally
on account of 'active=saved' mode.
Please perform an explicit 'copy running startup` to get ASCII configuration
```

Displaying FICON-Related Log Information

Example 28-25 and Example 28-26 display the logging information for FICON-related configurations.

Example 28-25 Displays Logging Levels for the FICON Feature

```
switch# show logging level ficon
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
ficon                   2                       2

0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Example 28-26 Displays FICON-Related Log File Contents

```
switch# show logging logfile
...
2004 Feb 25 15:38:50 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 13:22:04.
131183%$ Interface fc1/8 is up in mode F
 2004 Feb 25 15:38:50 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 13:22:04.
131217%$ Interface fc1/9 is up in mode F
...
2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
22:23.131121%$ Interface fc2/1, vsan 75 is up
 2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
22:23.131121%$ Interface fc2/2, vsan 75 is up
 2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
...
2004 Feb 25 23:22:36 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 21:05:42.
99916%$ Interface fc3/6 is up in mode F
 2004 Feb 25 23:22:37 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 21:05:43.
...
```

Default Settings

Table 28-3 lists the default settings for FICON features.

Table 28-3 Default FICON Settings

Parameters	Default
FICON feature	Disabled.
Port numbers	Same as port addresses.
FC ID last byte value	0 (zero).
EBCDIC format option	US-Canada.
Switch offline state	Hosts are allowed to move the switch to an offline state.
Mainframe users	Allowed to configure FICON parameters on Cisco MDS switches.
Clock in each VSAN	Same as the switch hardware clock.

Table 28-3 Default FICON Settings (continued)

Parameters	Default
Host clock control	Allows host to set the clock on this switch.
SNMP users	Configure FICON parameters.
Port address	Not blocked
Prohibited ports	Ports90–253 and 255 for the Cisco MDS 9200 Series switches. Ports250–253 and 255 for the Cisco MDS 9500 Series switches.

CHAPTER 29

Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

Common Information Model

Common Information Model (CIM) is an object-oriented information model that extends the existing standards for describing management information in a network/enterprise environment.

This section contains the following sections:

About CIM

CIM messages are independent of platform and implementation because they are encoded in N Extensible Markup Language (XML). CIM consists of a specification and a schema. The specification defines the syntax and rules for describing management data and integrating with other management models. The schema provides the actual model descriptions for systems, applications, networks, and devices.

For more information about CIM, refer to the specification available through the Distributed Management Task Force (DMTF) website at the following URL: http://www.dmtf.org/

For further information about Cisco MDS 9000 Family support for CIM servers, refer to the *Cisco MDS 9000 Family CIM Programming Reference Guide*.

A CIM client is required to access the CIM server. The client can be any client that supports CIM.

Configuring Added Security on a CIM Server

For added security, you can install an SSL certificate to encrypt the login information and enable the HTTPS server before enabling the CIM server. The CIM server is disabled by default. If you do not enable the HTTPS server, the standard HTTP server is enabled (default).

To configure a CIM server using the HTTPS protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# cimserver certificate bootflash:simserver.pem`	Installs a Secure Socket Layer (SSL) certificate specified in the file named with a .pem extension.
	`switch(config)# cimserver clearcertificate Certificate1`	Optional. Clears the specified SSL certificate (Certificate1).
Step 3	`switch(config)# cimserver enableHttps`	Enables HTTPS (secure protocol).
	`switch(config)# no cimserver enableHttps`	Optional. Disables HTTPS (default).
Step 4	`switch(config)# cimserver enable`	Enables the CIM server.
	`switch(config)# no cimserver enable`	Optional. Disables the CIM server (default).

To configure a CIM server using the HTTP protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# cimserver enable`	Enables the CIM server using the default HTTP (non-secure) protocol.
	`switch(config)# no cimserver enable`	Optional. Disables the CIM server (default).
	`switch(config)# no cimserver enableHttp`	Optional. Disables HTTP.
	`switch(config)# cimserver enableHttp`	Optional. Enables HTTP and reverts to the switch default.

Displaying CIM Information

To display CIM information, use the **show cimserver** command (see Example 29-1 through Example 29-4).

Example 29-1 Displays CIM Server Certificate Files

```
switch# show cimserver certificateName
cimserver certificate file name is servcert.pem
```

Example 29-2 Displays the CIM Server Configuration

```
switch# show cimserver
 cimserver is enabled
 cimserver Http is not enabled
 cimserver Https is enabled
 cimserver certificate file name is servcert.pem
```

Example 29-3 Displays the CIM Server HTTPS Status

```
switch# show cimserver httpsstatus
 cimserver Https is enabled
```

Example 29-4 Displays the CIM Server HTTP Status

```
switch# show cimserver httpstatus
 cimserver Http is not enabled
```

Fibre Channel Time Out Values

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

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

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

This section includes the following topics:

Timer Configuration Across All VSANs

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

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

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

To configure Fibre Channel timers across all VSANs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)`	Enters configuration mode.
Step 2	`switch(config)# fctimer R_A_TOV 6000`	Configures the R_A_TOV value for all VSANs to be 6000 msec. This type of configuration is not permitted unless all VSANs are suspended.

Timer Configuration Per-VSAN

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.

Caution You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.

Note This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.

To configure per-VSAN Fiber Channel timers, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)`	Enters configuration mode.

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

About fctimer Distribution

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

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

See Chapter 6, "Using the CFS Infrastructure," for more information on the CFS application.

Enabling or Disabling fctimer Distribution

To enable or disable fctimer fabric distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fctimer distribute`	Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.
	`switch(config)# no fctimer distribute`	Disables (default) fctimer configuration distribution to all switches in the fabric.

Committing fctimer Changes

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

To commit the fctimer configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fctimer commit`	Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database.

Discarding fctimer Changes

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

To discard the fctimer configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fctimer abort`	Discards the fctimer configuration changes in the pending database and releases the fabric lock.

Fabric Lock Override

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

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

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

```
switch# clear fctimer session
```

Database Merge Guidelines

See the "CFS Merge Support" section on page 6-8 for detailed concepts.

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
  - The merge protocol is not implemented for distribution of the fctimer values—you must manually merge the fctimer values when a fabric is merged.The per-VSAN fctimer configuration is distributed in the physical fabric.
  - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
  - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.

Note The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or abort the pending configurations before performing any more operations.

Displaying Configured fctimer Values

Use the **show fctimer** command to display the configured fctimer values (see Examples 29-5 and 29-6).

Example 29-5 Displays Configured Global TOVs

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

Note The F_S_TOV constant, though not configured, is displayed in the output of the **show fctimer** command.

Example 29-6 Displays Configured TOVs for a Specified VSAN

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

World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see Table 29-1).

Table 29-1 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.

This section includes the following topics:

- Displaying WWN Information, page 29-7
- Link Initialization WWN Usage, page 29-8
- Configuring a Secondary MAC Address, page 29-8

Displaying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. See Examples 29-7 to 29-9.

Example 29-7 Displays the Status of All WWNs

```
switch# show wwn status
         Type 1 WWNs: Configured:     64 Available:     48 (75%) Resvd.: 16
    Types 2 & 5 WWNs: Configured: 524288 Available: 450560 (85%) Resvd.: 73728
NKAU & NKCR WWN Blks: Configured:   1760 Available:   1760 (100%)
        Alarm Status:      Type1:   NONE Types 2&5:   NONE
```

Example 29-8 Displays Specified Block ID Information

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

Example 29-9 Displays the WWN for a Specific Switch

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

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco SAN-OS software release:

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

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

Note As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

Configuring a Secondary MAC Address

To allocate secondary MAC addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# wwn secondary-mac 00:99:55:77:55:55 range 64`
	`This command CANNOT be undone.`
	`Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55`
	`Please enter the mac address RANGE again: 64`
	`From now on WWN allocation would be based on new MACs.`
	`Are you sure? (yes/no) no`
	`You entered: no. Secondary MAC NOT programmed`
		Configures the secondary MAC address. This command cannot be undone.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs that do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b) (see the "FC ID Allocation for HBAs" section on page 29-8).

To allow further scalability for switches with numerous ports, the Cisco SAN-OS software maintains a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also know as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. Hence 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. Irrespective of the kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

This section includes the following topics:

Default Company ID list

All switches in the Cisco MDS 9000 Family that ship with Cisco SAN-OS Release 2.0(1b) or later, 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. Hence even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.

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

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

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

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

To allocate company IDs, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcid-allocation area company-id 0x003223`	Adds a new company ID to the default list.
	`switch(config)# no fcid-allocation area company-id 0x00E069`	Deletes a company ID from the default list.
	`switch(config)# fcid-allocation area company-id 0x003223`	Adds a new company ID to the default list.

Verifying the Company ID Configuration

You can view the configured company IDs by issuing the **show fcid-allocation area** command (see Example 29-10). 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.

Example 29-10 Displays the List of Default and Configured Company IDs

```
switch# show fcid-allocation area
FCID area allocation company id info:
    00:50:2E <---------------------- Default entry
    00:50:8B
    00:60:B0
    00:A0:B8
    00:E0:69
    00:30:AE + <------------------- User-added entry
    00:32:23 +

    00:E0:8B * <------------- Explicitly deleted entry (from the original default list)
Total company ids: 7
+ - Additional user configured company ids.
* - Explicitly deleted company ids from default list.
```

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

You can also view or obtain the company IDs in a specific WWN by issuing the **show fcid-allocation company-id-from-wwn** command (see Example 29-11). Some WWN formats do not support company IDs. In these cases, you many need to configure the FC ID persistent entry.

Example 29-11 Displays the Company ID for the Specified WWN

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

Switch Interoperability

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

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way, thus resulting in 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 more amiable standards compliant implementation.

Note For more information on configuring interoperability for the Cisco MDS 9000 Family switches, refer to the *Cisco MDS Family Switch-to-Switch Interoperability Configuration Guide*.

This section includes the following topics:

About Interop Mode

Cisco SAN-OS software supports the following four interop modes:

- Mode 1— Standards based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).
- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*.

Table 29-2 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Table 29-2 Changes in Switch Behavior 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 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior 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** Brocade uses the **cfgsave** command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN. **Note** Interop modes cannot be enabled on FICON-enabled VSANs.
TE ports and PortChannels	TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS 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.

Table 29-2 Changes in Switch Behavior When Interoperability Is Enabled (continued)

Switch Feature	Changes if Interoperability Is Enabled
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. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can be configured in **no interop** (default) mode or in any of the **interop** modes.

Configuring Interop Mode 1

The interop mode1 in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.

Note Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connecting from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

To configure interop mode 1 in any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.

```
switch# config t
switch(config)# vsan database
switch(config-vsan-db)# vsan 1 interop 1
switch(config-vsan-db)# exit
switch(config)#
```

Note You cannot enable interop modes on FICON-enabled VSANs.

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.

```
switch(config)# fcdomain domain 100 preferred vsan 1
```

In Cisco MDS 9000 switches, the default is to request an ID from the principal switch. If the preferred option is used, Cisco MDS 9000 switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the static option is used, the Cisco MDS 9000 switches do not join the fabric unless the principal switch agrees and assigns the requested ID.

Note When changing the domain ID, the FC IDs assigned to N ports also change.

Step 3 Change the Fibre Channel timers (if they have been changed from the system defaults).

Note The Cisco MDS 9000, 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 Cisco MDS domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.

```
switch(config)# fcdomain restart disruptive vsan 1
```

or

- Do not force a fabric reconfiguration.

```
switch(config# fcdomain restart vsan 1
```

Verifying Interoperating Status

This section highlights the commands used to verify if the fabric is up and running in interoperability mode.

To verify the resulting status of issuing the interoperability command in any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Use the **show version** command to verify the version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 1.0.8
  loader:    version 1.1(2)
  kickstart: version 2.0(1) [build 2.0(0.6)] [gdb]
  system:    version 2.0(1) [build 2.0(0.6)] [gdb]

  BIOS compile time:       08/07/03
  kickstart image file is: bootflash:///m9500-sf1ek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time:  10/25/2010 12:00:00
  system image file is:    bootflash:///m9500-sf1ek9-mzg.2.0.0.6.bin
  system compile time:     10/25/2020 12:00:00


Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:           0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
    Reason: Reset Requested by CLI command reload
    System version: 2.0(0.6)
    Service:
```

Step 2 Use the **show interface brief** command to verify if the interface states are as required by your configuration.

```
switch# show int brief
Interface  Vsan   Admin  Admin   Status          Oper  Oper   Port-channel
                  Mode   Trunk                   Mode  Speed
                         Mode                           (Gbps)
--------------------------------------------------------------------------
fc2/1      1      auto   on      up              E     2      --
fc2/2      1      auto   on      up              E     2      --
fc2/3      1      auto   on      fcotAbsent      --    --     --
fc2/4      1      auto   on      down            --    --     --
fc2/5      1      auto   on      down            --    --     --
fc2/6      1      auto   on      down            --    --     --
fc2/7      1      auto   on      up              E     1      --
fc2/8      1      auto   on      fcotAbsent      --    --     --
fc2/9      1      auto   on      down            --    --     --
fc2/10     1      auto   on      down            --    --     --
```

Step 3 Use the **show run** command to verify if you are running the desired configuration.

```
switch# show run
Building Configuration...

 interface fc2/1
no shutdown

 interface fc2/2
no shutdown

 interface fc2/3
 interface fc2/4
 interface fc2/5
 interface fc2/6
 interface fc2/7
no shutdown

 interface fc2/8
 interface fc2/9
 interface fc2/10

<snip>

interface fc2/32

 interface mgmt0
ip address 6.1.1.96 255.255.255.0
switchport encap default
no shutdown

vsan database
vsan 1 interop

boot system bootflash:/m9500-system-253e.bin sup-1
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-1
boot system bootflash:/m9500-system-253e.bin sup-2
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-2
callhome

fcdomain domain 100 preferred vsan 1

ip route 6.1.1.0 255.255.255.0 6.1.1.1
ip routing
```

```
line console
  databits 5
  speed 110
logging linecard
ssh key rsa 512 force
ssh server enable
switchname MDS9509
username admin password 5 $1$Li8/fBYX$SNc72.xt4nTXpSnR9OUFB/ role network-admin
```

Step 4 Use the **show vsan** command to verify if the interoperability mode is active.

```
switch# show vsan 1
vsan 1 information
         name:VSAN0001 stalactites
         interoperability mode:yes <---------- verify mode
         loadbalancing:src-id/dst-id/oxid
         operational state:up
```

Step 5 Use the **show fcdomain vsan** command to verify the domain ID.

```
switch# show fcdomain vsan 1
The local switch is a Subordinated Switch.

Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:05:30:00:51:1f
        Running fabric name: 10:00:00:60:69:22:32:91
        Running priority: 128
        Current domain ID: 0x64(100) <----------verify domain id

Local switch configuration information:
        State: Enabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 41:6e:64:69:61:6d:6f:21
        Configured priority: 128
        Configured domain ID: 0x64(100) (preferred)

Principal switch run time information:
        Running priority: 2

Interface               Role          RCF-reject
----------------    -------------    ------------
fc2/1               Downstream       Disabled
fc2/2               Downstream       Disabled
fc2/7               Upstream         Disabled
----------------    -------------    ------------
```

Step 6 Use the **show fcdomain domain-list vsan** command to verify the local principal switch status.

```
switch# show fcdomain domain-list vsan 1

Number of domains: 5
Domain ID              WWN
---------    -----------------------
 0x61(97)    10:00:00:60:69:50:0c:fe
 0x62(98)    20:01:00:05:30:00:47:9f
 0x63(99)    10:00:00:60:69:c0:0c:1d
0x64(100)    20:01:00:05:30:00:51:1f [Local]
0x65(101)    10:00:00:60:69:22:32:91 [Principal]
---------    -----------------------
```

Step 7 Use the **show fspf internal route vsan** command to verify the next hop and destination for the switch.

```
switch# show fspf internal route vsan 1
```

```
FSPF Unicast Routes
---------------------------
 VSAN Number  Dest Domain   Route Cost    Next hops
---------------------------------------------
           1     0x61(97)          500      fc2/2
           1     0x62(98)         1000      fc2/1
                                          fc2/2
           1     0x63(99)          500      fc2/1
           1    0x65(101)         1000      fc2/7
```

Step 8 Use the **show fcns data vsan** command to verify the name server information.

```
switch# show fcns data vsan 1
VSAN 1:
--------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------
0x610400    N     10:00:00:00:c9:24:3d:90 (Emulex)     scsi-fcp
0x6105dc    NL    21:00:00:20:37:28:31:6d (Seagate)    scsi-fcp
0x6105e0    NL    21:00:00:20:37:28:24:7b (Seagate)    scsi-fcp
0x6105e1    NL    21:00:00:20:37:28:22:ea (Seagate)    scsi-fcp
0x6105e2    NL    21:00:00:20:37:28:2e:65 (Seagate)    scsi-fcp
0x6105e4    NL    21:00:00:20:37:28:26:0d (Seagate)    scsi-fcp
0x630400    N     10:00:00:00:c9:24:3f:75 (Emulex)     scsi-fcp
0x630500    N     50:06:01:60:88:02:90:cb              scsi-fcp
0x6514e2    NL    21:00:00:20:37:a7:ca:b7 (Seagate)    scsi-fcp
0x6514e4    NL    21:00:00:20:37:a7:c7:e0 (Seagate)    scsi-fcp
0x6514e8    NL    21:00:00:20:37:a7:c7:df (Seagate)    scsi-fcp
0x651500    N     10:00:00:e0:69:f0:43:9f (JNI)

Total number of entries = 12
```

Note The Cisco MDS name server shows both local and remote entries, and does not time out the entries.

Default Settings

Table 29-4 lists the default settings for the features included in this chapter.

Table 29-3 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.

Table 29-3 Default Settings for Advanced Features (continued)

Parameters	Default
Local capture frame limit s	10 frames.
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.

Table 29-4 Default Settings for Advanced Features

Parameters	Default
D_S_TOV	5,000 msec
E_D_TOV	2,000 msec
R_A_TOV	10,000 msec
Interop mode	Disabled

PART 5

Security

CHAPTER 30

Configuring FIPS

The Federal Information Processing Standards (FIPS) Publication 140-2, *Security Requirements for Cryptographic Modules*, details the U.S. government requirements for cryptographic modules. FIPS 140-2 specifies that a cryptographic module shall be a set of hardware, software, firmware, or some combination thereof that implements cryptographic functions or processes, including cryptographic algorithms and, optionally, key generation, and is contained within a defined cryptographic boundary.

FIPS specifies certain crypto algorithms as secure, and it also identifies which algorithms should be used if a cryptographic module is to be called FIPS compliant.

Note Cisco MDS SAN-OS Release 3.1(1) implements FIPS features and is currently in the certification process with the U.S. government, but it is not FIPS compliant at this time.

This chapter includes the following sections:

Configuration Guidelines

Follow these guidelines before enabling FIPS mode.

- Make your passwords a minimum of eight characters in length.
- Disable Telnet. Users should log in using SSH only.
- Disable remote authentication through RADIUS/TACACS+. Only users local to the switch can be authenticated.
- Disable SNMP v1 and v2. Any existing user accounts on the switch that have been configured for SNMPv3 should be configured only with SHA for authentication and AES/3DES for privacy.
- Disable VRRP.
- Delete all IKE policies that either have MD5 for authentication or DES for encryption. Modify the policies so they use SHA for authentication and 3DES/AES for encryption.
- Delete all SSH Server RSA1 key-pairs.

Enabling FIPS Mode

To enable FIPS mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fips mode enable`	Enables FIPS mode.
	`switch(config)# no fips mode enable`	Disables FIPS mode.

Checking for FIPS Status

To view FIPS status, enter the **show fips status** command.

FIPS Self-Tests

A cryptographic module must perform power-up self-tests and conditional self-tests to ensure that it is functional.

Note FIPS power-up self-tests automatically run when FIPS mode is enabledby entering the **fips mode enable** command. A switch is in FIPS mode only after all self-tests are successfully completed. If any of the self-tests fail, then the switch is rebooted.

Power-up self-tests run immediately after FIPS mode is enabled. A cryptographic algorithm test using a known answer must be run for all cryptographic functions for each FIPS 140-2-approved cryptographic algorithm implemented on the Cisco MDS 9000 Family.

Using a known-answer test (KAT), a cryptographic algorithm is run on data for which the correct output is already known, and then the calculated output is compared to the previously generated output. If the calculated output does not equal the known answer, the known-answer test fails.

Conditional self-tests must be run when an applicable security function or operation is invoked. Unlike the power-up self-tests, conditional self-tests are executed each time their associated function is accessed.

Conditional self-tests include the following:

- Pair-wise consistency test—This test is run when a public-private key-pair is generated.
- Continuous random number generator test—This test is run when a random number is generated.

Both of these tests automatically run when a switch is in FIPS mode.

CHAPTER 31

Configuring Users and Common Roles

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use the CLI to modify a role that was created using SNMP and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

This chapter includes the following sections:

Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have permission to access that command.

This section includes the following topics:

About Roles

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to **debug** commands, then if Joe belongs to both role1 and role2, he can access configuration as well as **debug** commands.

Note If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.

Tip Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

Configuring Roles and Profiles

To create an additional role or to modify the profile for an existing role, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# role name techdocs` `switch(config-role)#`	Places you in the mode for the specified role (techdocs). **Note** The role submode prompt indicates that you are now in the role submode. This submode is now specific to the techdocs group.
	`switch(config)# no role name techdocs`	Deletes the role called techdocs.
Step 3	`switch(config-role)# description Entire Tech Docs group`	Assigns a description to the new role. The description is limited to one line and can contain spaces.
	`switch(config-role)# no description`	Resets the description for the Tech Docs group.

Note Only users belonging to the network-admin role can create roles.

Configuring Rules and Features for Each Role

Up to 16 rules can be configured for each role. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

For example, if user A is permitted to perform all **show** commands, user A cannot view the output of the **show role** command if user A does not belong to the network-admin role

The **rule** command specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, or interface).

Note In this case, **exec** commands refer to all commands in the EXEC mode that do not fall in the **show**, **debug**, and **clear** command categories.

Modifying Profiles

To modify the profile for an existing role, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# role name sangroup` `switch(config-role)#`	Places you in role configuration submode for the existing role sangroup.

Step 3	`switch(config-role)# rule 1 permit config`
	`switch(config-role)# rule 2 deny config feature fspf`
	`switch(config-role)# rule 3 permit debug feature zone`
	`switch(config-role)# rule 4 permit exec feature fcping`
		Allows users belonging to the sangroup role to perform all configuration commands except f**spf config** commands. They can also perform **zone debug** commands and the **fcping** EXEC mode command.

Step 4	`switch(config-role)# no rule 4`	Deletes rule 4, which no longer permits the sangroup to perform the **fcping** command.

In Step 3, rule 1 is applied first, thus permitting sangroup users access to all **config** commands. Rule 2 is applied next, denying FSPF configuration to sangroup users. As a result, sangroup users can perform all other **config** commands, except **fspf** configuration commands.

Note The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fspf** rule first and issued the **permit config** rule next, you would be allowing all sangroup users to perform all configuration commands because the second rule globally overrode the first rule.

Configuring the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license (see Chapter 3, "Obtaining and Installing Licenses").

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN policy for any role is permit, which allows tasks to be performed for all VSANs. You can configure a role that only allows tasks to be performed for a selected set of VSANs. To selectively allow VSANs for a role, set the VSAN policy to deny, and then set the configuration to permit or the appropriate VSANs.

Note Users configured in roles where the VSAN policy is set to deny cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.

Tip Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to deny are referred to as VSAN-restricted users.

Modifying the VSAN Policy

To modify the VSAN policy for an existing role, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# role name sangroup` `switch(config-role)#`	Places you in role configuration submode for the sangroup role.
Step 3	`switch(config)# vsan policy deny` `switch(config-role-vsan)`	Changes the VSAN policy of this role to **deny** and places you in a submode where VSANs can be selectively permitted.
	`switch(config-role)# no vsan policy deny`	Deletes the configured VSAN role policy and reverts to the factory default (**permit**).
Step 4	`switch(config-role-vsan)# permit vsan 10-30`	Permits this role to perform the allowed commands for VSANs 10 through 30.
	`switch(config-role-vsan)# no permit vsan 15-20`	Removes the permission for this role to perform commands for VSANs 15 to 20. So, the role is now permitted to perform commands for VSAN 10 to 14, and 21 to 30.

Role Distributions

Role-based configurations use the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and to provide a single point of configuration for the entire fabric (see Chapter 6, "Using the CFS Infrastructure").

The following configurations are distributed:

- Role names and descriptions
- List of rules for the roles
- VSAN policy and the list of permitted VSANs

This section includes the following topics:

About Role Databases

Role-based configurations use two databases to accept and implement configurations.

- Configuration database—The database currently enforced by the fabric.
- Pending database—Your subsequent configuration changes are stored in the pending database. If you modify the configuration, you need to commit or discard the pending database changes to the configuration database. The fabric remains locked during this period. Changes to the pending database are not reflected in the configuration database until you commit the changes.

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the entire fabric. 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 change.

Committing Role-Based Configuration Changes

If you commit the changes made to the pending 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. The configuration database now contains the committed changes and the pending database is now cleared.

To commit role-based configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# role commit vsan 3`	Commits the role-based configuration changes.

Discarding Role-Based Configuration Changes

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

To discard role-based configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# role abort`	Discards the role-based configuration changes and clears the pending configuration database.

Enabling Role-Based Configuration Distribution

To enable role-based configuration distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# role distribute`	Enables role-based configuration distribution.
	`switch(config)# no role distribute`	Disables role-based configuration distribution (default).

Clearing Sessions

To forcibly clear the existing role session in the fabric, issue the **clear role session** command from any switch that is part of the initiated session.

Caution Any changes in the pending database are lost when you issue this command.

```
switch# clear role session
```

Database Merge Guidelines

Fabric merge does not modify the role database on a switch. If two fabrics merge, and the fabrics have different role databases, the software generates an alert message.

See the "CFS Merge Support" section on page 6-8 for detailed concepts.

- Verify that the role database is identical on all switches in the entire fabric.
- Be sure to edit the role database on any switch to the desired database and then commit it. This synchronizes the role databases on all the switches in the fabric.

Displaying Role-Based Information

Use the **show role** command to display rules configured on the switch. The rules are displayed by rule number and are based on each role. All roles are displayed if the role name is not specified. See Example 31-1.

Example 31-1 Displays Information for All Roles

```
switch# show role
Role: network-admin
Description: Predefined Network Admin group. This role cannot be modified
Access to all the switch commands

Role: network-operator
Description: Predefined Network Operator group. This role cannot be modified
Access to Show commands and selected Exec commands

Role: svc-admin
Description: Predefined SVC Admin group. This role cannot be modified
Access to all SAN Volume Controller commands

Role: svc-operator
Description: Predefined SVC Operator group. This role cannot be modified
Access to selected SAN Volume Controller commands

Role: TechDocs
  vsan policy: permit (default)

Role: sangroup
  Description: SAN management group
  vsan policy: deny
  Permitted vsans: 10-30

  ----------------------------------------------
  Rule    Type    Command-type      Feature
  ----------------------------------------------
    1.   permit    config                  *
    2.     deny    config               fspf
    3.   permit     debug               zone
    4.   permit      exec             fcping
```

Displaying Roles When Distribution is Enabled

Use the **show role** command to display the configuration database.

Use the **show role status** command to display whether distribution is enabled for role configuration, the current fabric status (locked or unlocked), and the last operation performed. See Example 31-2.

Example 31-2 Displays the Role Status Information

```
switch# show role status
Distribution: Enabled
Session State: Locked

Last operation (initiated from this switch): Distribution enable
Last operation status: Success
```

Use the **show role pending** command to display the pending role database.

Example 31-3 displays the output of the **show role pending** command by following this procedure:

1. Create the role called `myrole` using the **role name myrole** command.
2. Issue the **rule 1 permit config feature fspf** command.
3. Issue the **show role pending** command to see the output.

Example 31-3 Displays Information on the Pending Roles Database

```
switch# show role pending
Role: network-admin
Description: Predefined Network Admin group. This role cannot be modified
Access to all the switch commands

Role: network-operator
Description: Predefined Network Operator group. This role cannot be modified
Access to Show commands and selected Exec commands

Role: svc-admin
Description: Predefined SVC Admin group. This role cannot be modified
Access to all SAN Volume Controller commands

Role: svc-operator
Description: Predefined SVC Operator group. This role cannot be modified
Access to selected SAN Volume Controller commands

Role: TechDocs
  vsan policy: permit (default)

Role: sangroup
  Description: SAN management group
  vsan policy: deny
  Permitted vsans: 10-30

  ----------------------------------------------
  Rule    Type    Command-type      Feature
  ----------------------------------------------
    1.   permit    config                  *
    2.     deny    config               fspf
    3.   permit     debug               zone
    4.   permit      exec             fcping
```

```
Role: myrole
  vsan policy: permit (default)
  ----------------------------------------------
  Rule     Type     Command-type      Feature
  ----------------------------------------------
    1.   permit     config                fspf
```

Use the **show role pending-diff** command to display the differences between the pending and configuration role database. See Example 31-4.

Example 31-4 Displays the Differences Between the Two Databases

```
switch# show role pending-diff
    +Role: myrole
    +  vsan policy: permit (default)
    +  ----------------------------------------------
    +  Rule     Type     Command-type      Feature
    +  ----------------------------------------------
    +    1.   permit     config                fspf
```

Configuring Common Roles

The CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using the CLI and vice versa (see Figure 31-1).

Figure 31-1 Common Roles

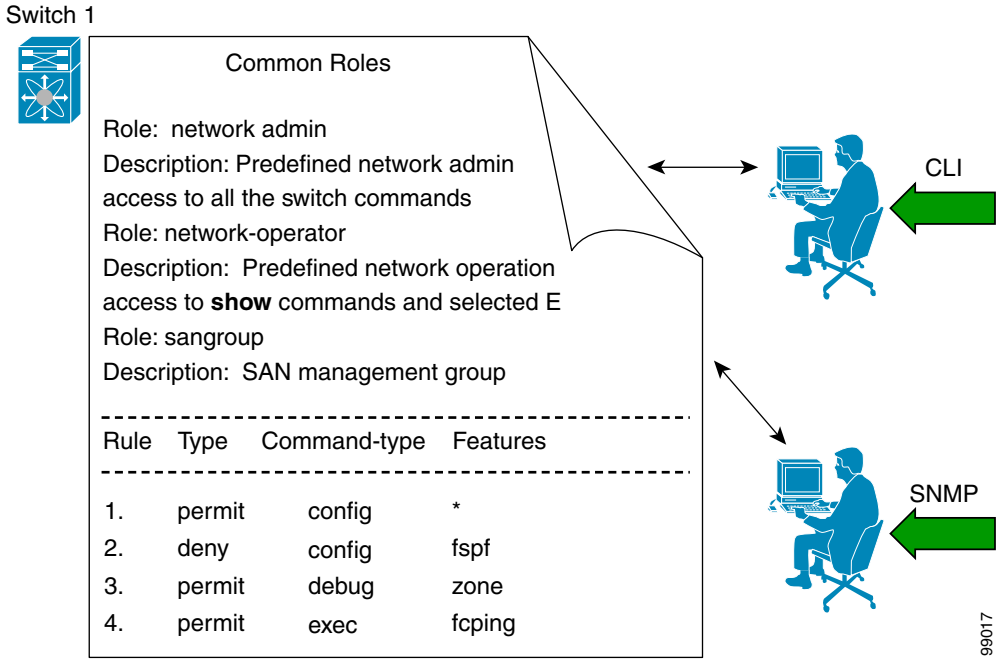

Each role in SNMP is the same as a role created or modified through the CLI (see the "Role-Based Authorization" section on page 31-1).

Each role can be restricted to one or more VSANs as required.

You can create new roles or modify existing roles using SNMP or the CLI.

- SNMP—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.
- CLI—Use the **role name** command.

Mapping of CLI Operations to SNMP

SNMP has only three possible operations: GET, SET, and NOTIFY. The CLI has five possible operations: DEBUG, SHOW, CONFIG, CLEAR, and EXEC.

Note NOTIFY does not have any restrictions like the syslog messages in the CLI.

Table 31-1 explains how the CLI operations are mapped to the SNMP operations.

Table 31-1 CLI Operation to SNMP Operation Mapping

CLI Operation	SNMP Operation
DEBUG	Ignored
SHOW	GET
CONFIG	SET
CLEAR	SET
EXEC	SET

Example 31-5 shows the privileges and rules mapping CLI operations to SNMP operations for a role named my_role.

Example 31-5 Displays CLI Operation to SNMP Operation Mapping

```
switch# show role name my_role
Role:my_role
  vsan policy:permit (default)
  ----------------------------------------------
  Rule    Type    Command-type      Feature
  ----------------------------------------------
    1.   permit     clear                 *
    2.     deny     clear               ntp
    3.   permit    config                 *
    4.     deny    config               ntp
    5.   permit     debug                 *
    6.     deny     debug               ntp
    7.   permit      show                 *
    8.     deny      show               ntp
    9.   permit      exec                 *
```

Note Although CONFIG is denied for NTP in rule 4, rule 9 allows the SET to NTP MIB objects because EXEC also maps to the SNMP SET operation.

Configuring User Accounts

Every Cisco MDS 9000 Family switch user has the account information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile.

The tasks explained in this section enable you to create users and modify the profile of an existing user. These tasks are restricted to privileged users as determined by your administrator.

This section includes the following topics:

About Users

The passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized (see the "SNMPv3 CLI User Management and AAA Integration" section on page 32-3).

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.

Note You can configure up to a maximum of 256 users on a switch.

Tip The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nscd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.

Note User passwords are not displayed in the switch configuration file.

Tip If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. "admin" is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.

Caution Cisco MDS SAN-OS does not support all numeric user names, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in.

Tip To issue commands with the **internal** keyword for troubleshooting purposes, you must have an account that is a member of the network-admin group.

Characteristics of Strong Passwords

A strong password has the following characteristics:

- At least eight characters long
- Does not contain many consecutive characters (such as “abcd”)
- Does not contain many repeating characters (such as “aaabbb”)
- Does not contain dictionary words
- Does not contain proper names
- Contains both upper- and lower-case characters
- Contains numbers

The following are examples of strong passwords:

- If2CoM18
- 2004AsdfLkj30
- Cb1955S21

Note Clear text passwords can only contain alphanumeric characters. The dollar sign ($) is not allowed.

Configuring Users

To configure a new user or to modify the profile of an existing user, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# username usam password abcd123AAA expire 2003-05-31`	Creates or updates the user account (usam) along with a password (abcd123AAA) that is set to expire on 2003-05-31. The password is limited to 64 characters. **Note** User account names must contain non-numeric characters.
	`switch(config)# username msam password 0 abcd12AAA role network-operator`	Creates or updates the user account (msam) along with a password (abcd12AAA) specified in clear text (indicated by 0). The password is limited to 64 characters. **Note** User account names must contain non-numeric characters.
	`switch(config)# username user1 password 5 !@*asdsfsdfjh!@df`	Specifies an encrypted (specified by 5) password (!@*asdsfsdfjh!@df) for the user account (user1).
Step 3	`switch(config)# username usam role network-admin`	Adds the specified user (usam) to the network-admin role.
	`switch(config)# no username usam role vsan-admin`	Deletes the specified user (usam) from the vsan-admin role.
Step 4	`switch(config)# username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSI YZ0EOdJ3l5RONWcwSgAuTUSrLk 3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8d veqts/8XQhqkNAFeGy4u8TJ2Us oreCU6DlibwkpzDafzKTpA5vB6FmHd2TI6Gnse9FUgKD5fs=`	Specifies the SSH key for an existing user account (admin).
	`switch(config)# no username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSI YZ0EOdJ3l5RONWcwSgAuTUSrLk 3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8d veqts/8XQhqkNAFeGy4u8TJ2Us oreCU6DlibwkpzDafzKTpA5vB6FmHd2TI6Gnse9FUgKD5fs=`	Deletes the SSH key for the user account (admin).
Step 5	`switch(config)# username usam ssh-cert-dn usam-dn dsa`	Specifies an SSH X.509 certificate distinguished name and DSA algorithm to use for authentication for an existing user account (usam).
	`switch(config)# username user1 ssh-cert-dn user1-dn rsa`	Specifies an SSH X.509 certificate distinguished name and RSA algorithm to use for authentication for an existing user account (user1).
	`switch(config)# no username admin ssh-cert-dn admin-dn dsa`	Removes the SSH X.509 certificate distinguished name for the user account (admin).

Logging Out Users

To log out another user on the switch, use the **clear user** command.

In the following example, the user named vsam is logged out from the switch.

```
switch# clear user vsam
```

Use the **show users** command to view a list of the logged in users (see Example 31-6).

Example 31-6 Displays All Logged in Users

```
switch# show users
admin    pts/7        Jan 12 20:56 (10.77.202.149)
admin    pts/9        Jan 12 23:29 (user.example.com)
admin    pts/10       Jan 13 03:05 (dhcp-10-10-1-1.example.com)
admin    pts/11       Jan 13 01:53 (dhcp-10-10-2-2.example.com)
```

Displaying User Account Information

Use the **show user-account** command to display configured information about user accounts. See Examples 31-7 to 31-8.

Example 31-7 Displays Information for a Specified User

```
switch# show user-account user1
user:user1
        this user account has no expiry date
        roles:network-operator
no password set. Local login not allowed
Remote login through RADIUS is possible
```

Example 31-8 Displays Information for All Users

```
switch# show user-account
show user-account
user:admin
        this user account has no expiry date
        roles:network-admin
user:usam
        expires on Sat May 31 00:00:00 2003
        roles:network-admin network-operator
user:msam
        this user account has no expiry date
        roles:network-operator
user:user1
        this user account has no expiry date
        roles:network-operator
no password set. local login not allowed
Remote login through RADIUS is possible
```

Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a server key-pair (see the "Generating the SSH Server Key-Pair" section on page 31-15).

Use the **ssh key** command to generate a server key.

Caution If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

This section includes the following topics:

About SSH

SSH provides secure communications to the Cisco SAN-OS CLI. You can use SSH keys for the following SSH options:

- SSH1
- SSH2, using RSA
- SSH2 using DSA

Generating the SSH Server Key-Pair

Be sure to have an SSH server key-pair with the appropriate version before enabling the SSH service. Generate the SSH server key-pair according to the SSH client version used. The number of bits specified for each key-pair ranges from 768 to 2048.

The SSH service accepts three types of key-pairs for use by SSH versions 1 and 2.

- The **rsa1** option generates the RSA1 key-pair for the SSH version 1 protocol.
- The **dsa** option generates the DSA key-pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA key-pair for the SSH version 2 protocol.

Caution If you delete all of the SSH keys, you cannot start a new SSH session.

To generate the SSH server key-pair, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

Step 2	`switch(config)# ssh key rsa1 1024`
	`generating rsa1 key.....`
	`generated rsa1 key`
		Generates the RSA1 server key-pair.

	`switch(config)# ssh key dsa 1024`
	`generating dsa key.....`
	`generated dsa key`
		Generates the DSA server key-pair.

	`switch(config)# ssh key rsa 1024`
	`generating rsa key.....`
	`generated rsa key`
		Generates the RSA server key-pair.

	`switch(config)# no ssh key rsa 1024` `cleared RSA keys`	Clears the RSA server key-pair configuration.

Specifying the SSH Key

You can specify an SSH key to log in using the SSH client without being prompted for a password. You can specify the SSH key in three different formats:

- Open SSH format
- IETF SECSH format
- Public Key Certificate in PEM format

To specify or delete the SSH key in OpenSSH format for a specified user, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ0EOdJ3l5RONWcwSgAuTUSrLk3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8dveqts/8XQhqkNAFeGy4u8TJ2UsoreCU6DlibwkpzDafzKTpA5vB6FmHd2TI6Gnse9FUgKD5fs=`	Specifies the SSH key for the user account (admin).
	`switch(config)# no username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ0EOdJ3l5RONWcwSgAuTUSrLk3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8dveqts/8XQhqkNAFeGy4u8TJ2UsoreCU6DlibwkpzDafzKTpA5vB6FmHd2TI6Gnse9FUgKD5fs=`	Deletes the SSH key for the user account (admin).

To specify or delete the SSH key in IETF SECSH format for a specified user, follow these steps:

	Command	Purpose
Step 1	`switch# copy tftp://10.10.1.1/secsh_file.pub bootflash:secsh_file.pub`	Downloads the file containing the SSH key in IETF SECSH format.
Step 2	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 3	`switch(config)# username admin sshkey file bootflash:secsh_file.pub`	Specifies the SSH key for the user account (admin).
	`switch(config)# no username admin sshkey file bootflash:secsh_file.pub`	Deletes the SSH key for the user account (admin).

To specify or delete the SSH key in PEM-formatted Public Key Certificate form for a specified user, follow these steps:

	Command	Purpose
Step 1	`switch# copy tftp://10.10.1.1/cert.pem bootflash:cert.pem`	Downloads the file containing the SSH key in PEM-formatted Public Key Certificate form.
Step 2	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 3	`switch(config)# username admin sshkey file bootflash:cert.pem`	Specifies the SSH key for the user account (usam).
	`switch(config)# no username admin sshkey file bootflash:cert.pem`	Deletes the SSH key for the user account (usam).

Overwriting a Generated Key-Pair

If the SSH key-pair option is already generated for the required version, you can force the switch to overwrite the previously generated key-pair.

To overwrite the previously generated key-pair, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

Step 2	`switch(config)# ssh key dsa 768`
	`ssh key dsa 512`
	`dsa keys already present, use force`
	`option to overwrite them`
	`switch(config)# ssh key dsa 512 force`
	`deleting old dsa key.....`
	`generating dsa key.....`
	`generated dsa key`
		Tries to set the server key-pair. If a required server key-pair is already configured, use the **force** option to overwrite that server key-pair.
		Deletes the old DSA key and sets the server key-pair using the new bit specification.

Clearing SSH Hosts

The **clear ssh hosts** command clears the existing list of trusted SSH hosts and reallows you to use SCP/SFTP along with the **copy** command for particular hosts.

When you use SCP/SFTP along with the **copy** command, a list of trusted SSH hosts are built and stored within the switch (see Example 31-9).

Example 31-9 Using SCP/SFTP to Copy Files

```
switch# copy scp://abcd@10.10.1.1/users/abcd/abc
bootflash:abc The authenticity of host '10.10.1.1 (10.10.1.1)'
can't be established.
RSA1 key fingerprint is 01:29:62:16:33:ff:f7:dc:cc:af:aa:20:f8:20:a2:db.
Are you sure you want to continue connecting (yes/no)? yes
Added the host to the list of known hosts
(/var/home/admin/.ssh/known_hosts). [SSH key information about the host is
stored on the switch]
abcd@10.10.1.1's password:
switch#
```

If a host's SSH key changes before you use SCP/SFTP along with the **copy** command, you will receive an error (see Example 31-10).

Example 31-10 Using SCP/SFTP to Copy Files—Error Caused by SSH Key Change

```
switch# copy scp://apn@10.10.1.1/isan-104
bootflash:isan-ram-1.0.4
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@    WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!     @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA1 host key has just been changed.
The fingerprint for the RSA1 key sent by the remote host is
36:96:ca:d7:29:99:79:74:aa:4d:97:49:81:fb:23:2f.
Please contact your system administrator.
Add correct host key in /mnt/pss/.ssh/known_hosts to get rid of this
message.
Offending key in /mnt/pss/.ssh/known_hosts:2
RSA1 host key for 10.10.1.1 has changed and you have requested strict
checking.
```

Enabling SSH or Telnet Service

By default, the SSH service is disabled.

To enable or disable the SSH service, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ssh server enable` `updated`	Enables the use of the SSH service.
	`switch(config)# no ssh server enable` `updated`	Disables (default) the use of the SSH service and resets the switch to its factory defaults.

Displaying SSH Protocol Status

Use the **show ssh server** command to display the status of the SSH protocol (enabled or disabled) and the versions that are enabled for that switch (see Example 31-11).

Example 31-11 Displays SSH Protocol Status

```
switch# show ssh server
ssh is enabled
version 1 enabled
version 2 enabled
```

Use the **show ssh key** command to display the server key-pair details for the specified key or for all keys, (see Example 31-12).

Example 31-12 Displays Server Key-Pair Details

```
switch# show ssh key
rsa1 Keys generated:Sun Jan 13 07:16:26 1980
1024 35
fingerprint:
1024 67:76:02:bd:3e:8d:f5:ad:59:5a:1e:c4:5e:44:03:07
could not retrieve rsa key information
dsa Keys generated:Sun Jan 13 07:40:08 1980
ssh-dss
AAAAB3NzaC1kc3MAAABBAJTCRQOydNRel2v7uiO6Fix+OTn8eGdnnDVxw5eJs5OcOEXOyjaWcMMYsEgxc9ada1NElp
8Wy7GPMWGOQYj9CU0AAAAVAMCcWhNN18zFNOIPo7cU3t7d0iEbAAAAQBdQ8UAOi/Cti84qFb3kTqXlS9mEhdQUo0lH
cH5bw5PKfj2Y/dLR437zCBKXetPj4p7mhQ6Fq5os8RZtJEyOsNsAAABAA0oxZbPyWeR5NHATXiyXdPI7j9i8fgyn9F
NipMkOF2Mn75Mi/lqQ4NIq0gQNvQOx27uCeQlRts/QwI4q68/eaw=
fingerprint:
512 f7:cc:90:3d:f5:8a:a9:ca:48:76:9f:f8:6e:71:d4:ae
```

Note If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** CLI command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

SSH Authentication Using Digital Certificates

SSH authentication on the Cisco MDS 9000 Family switches provide X.509 digital certificate support for host authentication. An X.509 digital certificate is a data item that vouches for the origin and integrity of a message. It contains encryption keys for secured communications and is "signed" by a trusted certification authority (CA) to verify the identity of the presenter. The X.509 digital certificate support provides either DSA or RSA algorithms for authentication.

The certificate infrastructure uses the first certificate that supports the Secure Socket Layer (SSL) and is returned by the security infrastructure, either through query or notification. Verification of certificates is successful if the certificates are from any of the trusted CAs.

You can configure your switch for either SSH authentication using an X.509 certificate or SSH authentication using a Public Key Certificate, but not both. If either of them is configured and the authentication fails, you will be prompted for a password.

For more information on CAs and digital certificates, see Chapter 35, "Configuring Certificate Authorities and Digital Certificates."

Recovering the Administrator Password

You can recover the administrator password using one of two methods:

- From the CLI with a user name that has network-admin privileges.
- Power cycling the switch.

The following topics included in this section:

Using the CLI with Network-Admin Privileges

If you are logged in to, or can log into, switch with a user name that has network-admin privileges and then recover the administrator password, follow these steps:

Step 1 Use the **show user-accounts** command to verify that your user name has network-admin privileges.

```
switch# show user-account
user:admin
        this user account has no expiry date
        roles:network-admin

user:dbgusr
        this user account has no expiry date
        roles:network-admin network-operator
```

Step 2 If your user name has network-admin privileges, issue the **username** command to assign a new administrator password.

```
switch# config t
switch(config)# username admin password <new password>
switch(config)# exit
switch#
```

Step 3 Save the software configuration.

```
switch# copy running-config startup-config
```

Power Cycling the Switch

If you cannot start a session on the switch that has network-admin privileges, you must recover the administrator password by power cycling the switch.

Caution This procedure disrupts all traffic on the switch. All connections to the switch will be lost for 2 to 3 minutes.

Note You cannot recover the administrator password from a Telnet or SSH session. You must have access to the local console connection. See the "Starting a Switch in the Cisco MDS 9000 Family" section on page 5-2 for information on setting up the console connection.

To recover a administrator password by power cycling the switch, follow these steps:

Step 1 For Cisco MDS 9500 Series switches with two supervisor modules, remove the supervisor module in slot 6 from the chassis.

Note On the Cisco MDS 9500 Series, the password recovery procedure must be performed on the active supervisor module. Removing the supervisor module in slot 6 ensures that a switchover will not occur during the password recovery procedure.

Step 2 Power cycle the switch.

Step 3 Press the **Ctrl-]** key sequence when the switch begins its Cisco SAN-OS software boot sequence to enter the `switch(boot)#` prompt mode.

```
Ctrl-]
switch(boot)#
```

Step 4 Change to configuration mode.

```
switch(boot)# config terminal
```

Step 5 Issue the **admin-password** command to reset the administrator password.

```
switch(boot-config)# admin-password <new password>
```

For information on strong passwords, see the "Characteristics of Strong Passwords" section on page 31-12.

Step 6 Exit to the EXEC mode.

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

Step 7 Issue the **load** command to load the Cisco SAN-OS software.

```
switch(boot)# load bootflash:m9500-sf1ek9-mz.2.1.1a.bin
```

Caution If you boot a system image that is older than the image you used to store the configuration and do not use the **install all** command to boot the system, the switch erases the binary configuration and uses the ASCII configuration. When this occurs, you must use the **init system** command to recover your password.

Step 8 Log in to the switch using the new administrator password.

```
switch login: admin
Password: <new password>
```

Step 9 Reset the new password to ensure that is it is also the SNMP password for Fabric Manager.

```
switch# config t
switch(config)# username admin password <new password>
switch(config)# exit
switch#
```

Step 10 Save the software configuration.

```
switch# copy running-config startup-config
```

Step 11 Insert the previously removed supervisor module into slot 6 in the chassis.

Default Settings

Table 31-2 lists the default settings for all switch security features in any switch.

Table 31-2 Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS Switches	Network operator (network-operator).
AAA configuration services	Local.
Authentication port	1821.
Accounting port	1813.
Preshared key communication	Clear text.
RADIUS server time out	1 (one) second.
RADIUS server retries	Once.
TACACS+	Disabled.
TACACS+ servers	None configured.
TACACS+ server timeout	5 seconds.
AAA server distribution	Disabled.
VSAN policy for roles	Permit.
User account	No expiry (unless configured).

Table 31-2 Default Switch Security Settings (continued)

Parameters	Default
Password	None.
Accounting log size	250 KB.
SSH service	Disabled.
Telnet service	Enabled.

CHAPTER 32

Configuring SNMP

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using the CLI and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

This chapter includes the following sections:

About SNMP Security

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see Figure 32-1).

Figure 32-1 SNMP Security

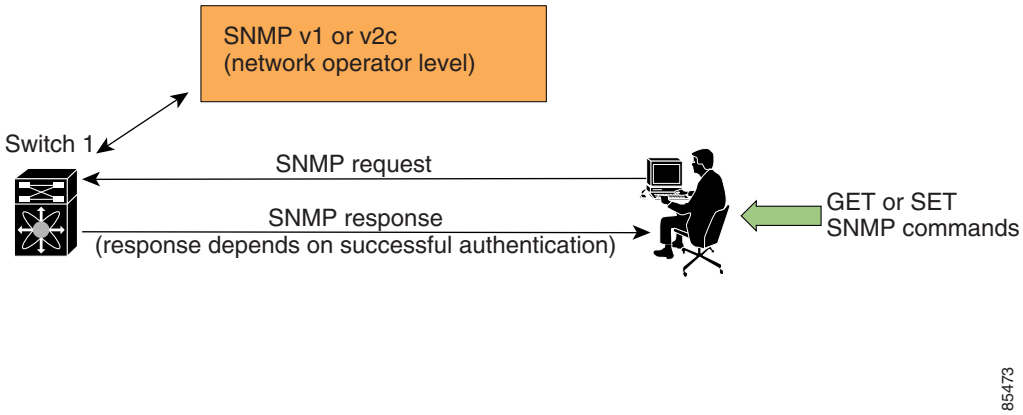

This section includes the following topics:

SNMP Version 1 and Version 2c

SNMP Version 1 (SNMPv1) and SNMP Version 2c (SNMPv2c) use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

SNMP Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Assigning SNMP Switch Contact and Location Information

You can assign the switch contact information, which is limited to 32 characters (without spaces) and the switch location.

To configure contact and location information, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# snmp-server contact NewUser`	Assigns the contact name for the switch.
	`switch(config)# no snmp-server contact NewUser`	Deletes the contact name for the switch.
Step 3	`switch(config)# snmp-server location SanJose`	Assigns the switch location.
	`switch(config)# no snmp-server location SanJose`	Deletes the switch location.

SNMPv3 CLI User Management and AAA Integration

The Cisco SAN-OS software implements RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

SNMPv3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

This section includes the following topics:

CLI and SNMP User Synchronization

Any configuration changes made to the user group, role, or password results in database synchronization for both SNMP and AAA.

To create an SNMP or CLI user, use either the **username** or **snmp-server user** commands.

- The `auth` passphrase specified in the **snmp-server user** command is synchronized as the password for the CLI user.
- The password specified in the **username** command is synchronized as the `auth` and `priv` passphrases for the SNMP user.

Users are synchronized as follows:

- Deleting a user using either command results in the user being deleted for both SNMP and the CLI.
- User-role mapping changes are synchronized in SNMP and the CLI.

Note When the passphrase/password is specified in localized key/encrypted format, the password is not synchronized.

- Existing SNMP users continue to retain the `auth` and `priv` passphrases without any changes.
- If the management station creates an SNMP user in the usmUserTable, the corresponding CLI user is created without any password (login is disabled) and will have the network-operator role.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs). SeeChapter 34, “Configuring IPv4 and IPv6 Access Control Lists,”

Group-Based SNMP Access

Note Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

Creating and Modifying Users

You can create users or modify existing users using or the CLI.

- SNMP—Create a user as a clone of an existing user in the usmUserTable on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC 2574.
- CLI—Create a user or modify an existing user using the **snmp-server user** command.

A network-operator and network-admin roles are available in a Cisco MDS 9000 Family switch. There is also a default-role if you want to use the GUI (Fabric Manager and Device Manager). You can also use any role that is configured in the Common Roles database (see the "User Accounts" section on page 37-10).

Tip All updates to the CLI security database and the SNMP user database are synchronized. You can use the SNMP password to log into either Fabric Manager or Device Manager. However, after you use the CLI password to log into Fabric Manager or Device Manager, you must use the CLI password for all future logins. If a user exists in both the SNMP database and the CLI database before upgrading to Cisco MDS SAN-OS Release 2.0(1b), then the set of roles assigned to the user becomes the union of both sets of roles after the upgrade.

This section includes the following topics:

About AES Encryption-Based Privacy

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm. The Cisco SAN-OS software uses AES as one of the privacy protocols for SNMP message encryption and conforms with RFC 3826.

The **priv** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The **priv** option along with the **aes-128** token indicates that this privacy password is for generating a 128-bit AES key.The AES priv password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.

Note For an SNMPv3 operation using the external AAA server, user configurations in the external AAA server require AES to be the privacy protocol to use SNMP PDU encryption.

Configuring SNMP Users from the CLI

The passphrase specified in the **snmp-server user** command and the **username** command are synchronized (see the "SNMPv3 CLI User Management and AAA Integration" section on page 32-3).

To create or modify SNMP users from the CLI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server user joe network-admin auth sha abcd1234`	Creates or modifies the settings for a user (joe) in the network-admin role using the HMAC-SHA-96 authentication password (abcd1234).
	`switch(config)# snmp-server user sam network-admin auth md5 abcdefgh`	Creates or modifies the settings for a user (sam) in the network-admin role using the HMAC-MD5-96 authentication password (abcdefgh).
	`switch(config)# snmp-server user Bill network-admin auth sha abcd1234 priv abcdefgh`	Creates or modifies the settings for a user (Bill) in the network-admin role using the HMAC-SHA-96 authentication level and privacy encryption parameters.
	`switch(config)# no snmp-server user usernameA`	Deletes the user (usernameA) and all associated parameters.
	`switch(config)# no snmp-server usam role vsan-admin`	Deletes the specified user (usam) from the vsan-admin role.
	`switch(config)# snmp-server user user1 network-admin auth md5 0xab0211gh priv 0x45abf342 localizedkey`	Specifies the password to be in localized key format (RFC 2574). The localized key is provided in hexadecimal format (for example, 0xacbdef).
	`switch(config)# snmp-server  user user2 auth md5 asdgfsadf priv aes-128 asgfsgkhkj`	Configures the user2 with the MD5 authentication protocol and AES-128 privacy protocol.
Step 3	`switch(config)# snmp-server user joe sangroup`	Adds the specified user (joe) to the sangroup role.
	`switch(config)# snmp-server user joe techdocs`	Adds the specified user (joe) to the techdocs role.

To create or modify passwords for SNMP users from the CLI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server user user1 role1 auth md5 0xab0211gh priv des 0x45abf342 localizedkey`	Specifies the password to be in localized key format using the DES option for security encryption.
	`switch(config)# snmp-server user user1 role2 auth sha 0xab0211gh priv aes-128 0x45abf342 localizedkey`	Specifies the password to be in localized key format using the 128-bit AES option for security encryption

Caution Avoid using the **localizedkey** option when configuring an SNMP user from the CLI. The localized keys are not portable across devices as they contain device engine ID information. If a configuration file is copied to the device, the passwords may not be set correctly if the configuration file was generated at a different device. Explicitly configure the desired passwords after copying the configuration into the device. Passwords specified with the **localizedkey** option are limited to 130 characters.

Note The **snmp-server user** command takes the engineID as an additional parameter. The engineID creates the notification target user (see the "Configuring the Notification Target User" section on page 32-12). If the engineID is not specified, the local user is created.

Enforcing SNMPv3 Message Encryption

By default the SNMP agent allows the securityLevel parameters of authNoPriv and authPriv for the SNMPv3 messages that use user-configured SNMPv3 message encryption with auth and priv keys.

To enforce the message encryption for a user, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server user testUser enforcePriv`	Enforces the message encryption for SNMPv3 messages using this user. **Note** You can only use this command for previously existing users configured with both auth and priv keys. When the user is configured to enforce privacy, for any SNMPv3 PDU request using securityLevel parameter of either noAuthNoPriv or authNoPriv, the SNMP agent responds with authorizationError.
	`switch(config)# no snmp-server user testUser enforcePriv`	Disables SNMPv3 message encryption enforcement.

Alternatively, you can enforce the SNMPv3 message encryption globally on all the users using the following commands:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server globalEnforcePriv`	Enforces the SNMPv3 message encryption for all the users on the switch.
	`switch(config)# no snmp-server globalEnforcePriv`	Disables global SNMPv3 message encryption enforcement.

Assigning SNMPv3 Users to Multiple Roles

The SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. After the initial SNMPv3 user creation, you can map additional roles for the user.

Note Only users belonging to a network-admin role can assign roles to other users.

To configure multiple roles for SNMPv3 users from the CLI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server user NewUser role1`	Creates or modifies the settings for an SNMPv3 user (NewUser) for the role1 role.
	`switch(config)# snmp-server user NewUser role2`	Creates or modifies the settings for an SNMPv3 user (NewUser) for the role2 role.
	`switch(config)# no snmp-server user User5 role2`	Removes role2 for the specified user (User5).

Adding or Deleting Communities

You can configure read-only or read-write access for SNMPv1 and SNMPv2 users. Refer to RFC 2576.

To create an SNMPv1 or SNMPv2c community, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# snmp-server community snmp_Community ro`	Adds read-only access for the specified SNMP community.
	`switch(config)# snmp-server community snmp_Community rw`	Adds read-write access for the specified SNMP community.
	`switch(config)# no snmp-server community snmp_Community`	Deletes access for the specified SNMP community (default).

SNMP Trap and Inform Notifications

You can configure the Cisco MDS switch to send notifications to SNMP managers when particular events occur.

Note Use the SNMP-TARGET-MIB to obtain more information on the destinations to which notifications are to be sent either as traps or as informs. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.

This section includes the following topics:

Tip The SNMPv1 option is not available with the **snmp-server host** *ip-address* **informs** command.

Configuring SNMPv2c Notifications

To configure SNMPv2c notifications using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server host 171.71.187.101 traps version 2c private udp-port 1163`	Configures the specified host to receive SNMPv2c traps using SNMPv2c community string (private).
	`switch(config)# no snmp-server host 171.71.187.101 traps version 2c private udp-port 2162`	Prevents the specified host from receiving SNMPv2c traps on the configured UDP port using SNMPv2c community string (private).
Step 3	`switch(config)# snmp-server host 171.71.187.101 informs version 2c private udp-port 1163`	Configures the specified host to receive SNMPv2c informs using SNMPv2c community string (private).
	`switch(config)# no snmp-server host 171.71.187.101 informs version 2c private udp-port 2162`	Prevents the specified host from receiving SNMPv2c informs on the configured UDP port using SNMPv2c community string (private).

To configure SNMPv2c notifications using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server host 2001:0DB8:800:200C::417A traps version 2c private udp-port 1163`	Configures the specified host to receive SNMPv2c traps using SNMPv2c community string (private).
	`switch(config)# no snmp-server host 2001:0DB8:800:200C::417A traps version 2c private udp-port 2162`	Prevents the specified host from receiving SNMPv2c traps on the configured UDP port using SNMPv2c community string (private).
Step 3	`switch(config)# snmp-server host 2001:0DB8:800:200C::417A informs version 2c private udp-port 1163`	Configures the specified host to receive SNMPv2c informs using SNMPv2c community string (private).
	`switch(config)# no snmp-server host 2001:0DB8:800:200C::417A informs version 2c private udp-port 2162`	Prevents the specified host from receiving SNMPv2c informs on the configured UDP port using SNMPv2c community string (private).

Note Switches can forward events (SNMP traps and informs) up to 10 destinations.

Configuring SNMPv3 Notifications

To configure SNMPv3 notifications using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server host 16.20.11.14 traps version 3 noauth testuser udp-port 1163`	Configures the specified host to receive SNMPv3 traps using SNMPv3 user (testuser) and securityLevel of noAuthNoPriv.
	`switch(config)# snmp-server host 16.20.11.14 informs version 3 auth testuser udp-port 1163`	Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthNoPriv.
	`switch(config)# snmp-server host 16.20.11.14 informs version 3 priv testuser udp-port 1163`	Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthPriv.
	`switch(config)# no snmp-server host 172.18.2.247 informs version 3 testuser noauth udp-port 2162`	Prevents the specified host from receiving SNMPv3 informs.

To configure SNMPv3 notifications using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server host 2001:0DB8:800:200C::417A traps version 3 noauth testuser udp-port 1163`	Configures the specified host to receive SNMPv3 traps using SNMPv3 user (testuser) and securityLevel of noAuthNoPriv.
	`switch(config)# snmp-server host 2001:0DB8:800:200C::417A informs version 3 auth testuser udp-port 1163`	Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthNoPriv.
	`switch(config)# snmp-server host 2001:0DB8:800:200C::417A informs version 3 priv testuser udp-port 1163`	Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthPriv.
	`switch(config)# no snmp-server host 2001:0DB8:800:200C::417A informs version 3 testuser noauth udp-port 2162`	Prevents the specified host from receiving SNMPv3 informs.

Note In the case of SNMPv3 notifications, the SNMP manager is expected to know the user credentials (authKey/PrivKey) based on the switch's engineID to authenticate and decrypt the SNMP messages.

Enabling SNMP Notifications

Notifications (traps and informs) are system alerts that the switch generates when certain events occur. You can enable or disable notifications. By default, no notification is defined or issued. If a notification name is not specified, all notifications are disabled or enabled.

Table 32-1 lists the CLI commands that enable the notifications for Cisco MDS MIBs.

Note The **snmp-server enable traps** CLI command enables both traps and informs, depending on how you configured . See notifications with the **snmp-server host** CLI command.

Table 32-1 Enabling SNMP Notifications

MIB	Related Commands
All notifications	**snmp-server enable traps**
CISCO-AAA-SERVER-MIB	**snmp-server enable traps aaa**
ENITY-MIB, CISCO-ENTITY-FRU-CONTROL-MIB, CISCO-ENTITY-SENSOR-MIB	**snmp-server enable traps entity** **snmp-server enable traps entity fru**
CISCO-FCC-MIB	**snmp-server enable traps fcc**
CISCO-DM-MIB	**snmp-server enable traps fcdomain**
CISCO-NS-MIB	**snmp-server enable traps fcns**
CISCO-FCS-MIB	**snmp-server enable traps fcs discovery-complete** **snmp-server enable traps fcs request-reject**

Table 32-1 Enabling SNMP Notifications (continued)

MIB	Related Commands
CISCO-FDMI-MIB	**snmp-server enable traps fdmi**
CISCO-FSPF-MIB	**snmp-server enable traps fspf**
CISCO-LICENSE-MGR-MIB	**snmp-server enable traps license**
IF-MIB	**snmp-server enable traps link**
CISCO-PSM-MIB	**snmp-server enable traps port-security**
CISCO-RSCN-MIB	**snmp-server enable traps rscn** **snmp-server enable traps rscn els** **snmp-server enable traps rscn ils**
SNMPv2-MIB	**snmp-server enable traps snmp** **snmp-server enable traps snmp authentication**
VRRP-MIB, CISCO-IETF-VRRP-MIB	**snmp-server enable traps vrrp**
CISCO-ZS-MIB	**snmp-server enable traps zone** **snmp-server enable traps zone default-zone-behavior-change** **snmp-server enable traps zone merge-failure** **snmp-server enable traps zone merge-success** **snmp-server enable traps zone request-reject** **snmp-server enable traps zone unsupp-mem**

The following notifications are enabled by default:

- **entity fru**
- **license**
- **link ietf-extended**

All other notifications are disabled by default.

To enable individual notifications, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server enable traps fcdomain`	Enables the specified SNMP (fcdomain) notification.
	`switch(config)# no snmp-server enable traps`	Disables the specified SNMP notification. If a notification name is not specified, all notifications are disabled.

You can use the **show snmp trap** command to display all the notifications and their status.

```
switch# show snmp trap
Trap type                              Enabled
---------                              -------
entity fru                              Yes
fcc                                     No
fcdomain                                No
fcns                                    No
fcs request-reject                      No
fcs discovery-complete                  No
```

```
fdmi                                  No
fspf                                  No
license                               Yes
port-security                         No
rscn els                              No
rscn ils                              No
snmp authentication                   No
vrrp                                  Yes
zone unsupported member               No
zone request-reject                   No
zone merge-failure                    No
zone merge-success                    No
zone default-zone-behavior-change     No
```

Configuring the Notification Target User

You must configure a notification target user on the switch for sending SNMPv3 inform notifications to the SNMP manager.

To configure the notification target user, use the following command:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03`	Configures the notification target user with the specified credentials for the SNMP manager with the specified engine ID.
	`switch(config)# no snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03`	Removes the notification target user.

The credentials of the notification target user are used for encrypting the SNMPv3 inform notification messages to the configured SNMP manager (as in the **snmp-server host** command).

Note For authenticating and decrypting the received INFORM PDU, the SNMP manager should have the same user credentials in its local configuration data store of users.

Configuring LinkUp/LinkDown Notifications for Switches

You can configure which linkUp/linkDown notifications to enable on switches. You can enable the following types of linkUp/linkDown notifications:

- Cisco—Only notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface.
- IETF—Only notifications (linkUp, linkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the notifications.

- IEFT extended—Only notifications (linkUp, linkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in the notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent. This is the default setting.
- IEFT Cisco—Only notifications (linkUp, linkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the linkUp and linkDown notifications.
- IEFT extended Cisco—Only notifications (linkUp, linkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in linkUp and linkDown notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent with the linkUp and linkDown notifications.

Note For more information on the varbinds defined in the IF-MIB specific to the Cisco Systems implementation, refer to the *Cisco MDS 9000 Family MIB Quick Reference.*

To configure the linkUp/linkDown notification for a switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 1	`switch(config)# snmp-server enable traps link`	Enables (default) only IETF extended linkUp/linkDown notifications.
	`switch(config)# snmp-server enable traps link cisco`	Enables only Cisco Systems defined cieLinkUp/cieLinkDown notifications.
	`switch(config)# snmp-server enable traps link ietf`	Enables only IETF linkUp/linkDown notifications.
	`switch(config)# snmp-server enable traps link ietf-extended`	Enables (default) only IETF extended linkUp/linkDown notifications with extra varbinds.
	`switch(config)# snmp-server enable traps link ietf cisco`	Enables IETF (linkUp/linkDown) and Cisco Systems defined (cieLinkUp/cieLinkDown) notifications.
	`switch(config)# snmp-server enable traps link ietf-extended cisco`	Enables IEFT (linkUp/linkDown) notifications with extra varbinds and Cisco Systems defined (cieLinkUp/cieLinkDown) notifications.
	`switch(config)# no snmp-server enable traps link`	Reverts to the default setting (IETF extended).

Configuring Up/Down SNMP Link-State Traps for Interfaces

By default, SNMP link-state traps are enabled for all interfaces. Whenever a link toggles its state from Up to Down or vice versa, an SNMP trap is generated.

In some instances, you may find that you have numerous switches with hundreds of interfaces, many of which do not require monitoring of the link state. In such cases, you may elect to disable link-state traps.

To disable SNMP link-state traps for specific interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface bay 6`	Specifies the interface on which to disable SNMP link-state traps.
	`switch(config-if)# no link-state-trap`	Disables SNMP link-state traps for the interface.
	`switch(config-if)# link-state-trap`	Enables SNMP link-state traps for the interface.

Whenever you disable an SNMP link-state trap for an interface, the command is also added to the running configuration of the system. To view the running configuration, enter the **show running-config** command for the interface.

```
switch# show running-config
version 3.1(2)
....
interface bay5
interface bay6
  no link-state-trap <-----command is added to the running configuration for the interface
interface bay7
...
```

To view the SNMP link-state trap configuration for a particular interface, enter the **show interface** command.

```
switch# show interface bay 6
bay6 is down (Administratively down)
    Hardware is Fibre Channel
    Port WWN is 20:0b:00:05:30:01:70:2c
    Admin port mode is auto, trunk mode is on
    snmp link-state traps are disabled

    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

Displaying SNMP Security Information

Use the **show snmp** commands to display configured SNMP information (see Example 32-1 and 32-6).

Example 32-1 Displays SNMP User Details

```
switch# show snmp user
_____________________________________________________________
                   SNMP USERS
_____________________________________________________________
```

```
User                          Auth  Priv(enforce) Groups
____                          ____  ______________ ______
admin                         md5   des(no)       network-admin

testusr                       md5   aes-128(no)   role111
                                                  role222
_____________________________________________________________________
 NOTIFICATION TARGET USERS (configured  for sending V3 Inform)
_____________________________________________________________________
User                          Auth  Priv
____                          ____  ____
testtargetusr                 md5   des
(EngineID 0:0:0:63:0:1:0:0:0:15:10:3)
```

Example 32-2 Displays SNMP Community Information

```
switch# show snmp community
Community                      Access
---------                      ------
private                         rw
public                          ro
v93RACqPNH                      ro
```

Example 32-3 Displays SNMP Host Information

```
switch# show snmp host
Host                           Port Version  Level  Type   SecName
____                           ____ _______  ______ ____   _______
171.16.126.34                  2162 v2c      noauth trap   public
171.16.75.106                  2162 v2c      noauth trap   public
...
171.31.58.97                   2162 v2c      auth   trap   public
...
```

The **show snmp** command displays counter information for SNMP contact, location, and packet settings. This command provides information that is used entirely by the Cisco MDS 9000 Family Fabric Manager (refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*). See Example 32-4.

Example 32-4 Displays SNMP Information

```
switch# show snmp
sys contact:
sys location:
1631 SNMP packets input
        0 Bad SNMP versions
        0 Unknown community name
        0 Illegal operation for community name supplied
        0 Encoding errors
        64294 Number of requested variables
        1 Number of altered variables
        1628 Get-request PDUs
        0 Get-next PDUs
        1 Set-request PDUs
152725 SNMP packets output
        0 Too big errors
        1 No such name errors
        0 Bad values errors
        0 General errors
Community                      Group / Access
---------                      ------
```

```
public                          rw

_____________________________________________________________________
                   SNMP USERS
_____________________________________________________________________

User                            Auth  Priv(enforce) Groups
____                            ____  _____________ ______
admin                           md5   des(no)       network-admin

testusr                         md5   aes-128(no)   role111
                                                    role222

_____________________________________________________________________
 NOTIFICATION TARGET USERS (configured  for sending V3 Inform)
_____________________________________________________________________

User                            Auth  Priv
____                            ____  ____
testtargetusr                   md5   des
(EngineID 0:0:0:63:0:1:0:0:0:15:10:3)
```

Example 32-5 Displays SNMP Engine IDs

```
switch# show snmp engineID
Local SNMP engineID: 800000090300053000851E
```

Example 32-6 Displays Information on SNMP Security Groups

```
switch# show snmp group
groupname: network-admin
security model: any
security level: noAuthNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active

groupname: network-admin
security model: any
security level: authNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active

groupname: network-operator
security model: any
security level: noAuthNoPriv
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active

groupname: network-operator
security model: any
security level: authNoPriv
```

```
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active
```

Default Settings

Table 32-2 lists the default settings for all SNMP features in any switch.

Table 32-2 Default SNMP Settings

Parameters	Default
User account	No expiry (unless configured).
Password	None.

CHAPTER 33

Configuring RADIUS and TACACS+

The authentication, authorization, and accounting (AAA) feature verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use RADIUS and TACACS+ protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using a AAA server. A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA servers or for only a specific AAA server. This security feature provides a central management capability for AAA servers.

This chapter includes the following sections:

Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family provides security to all management access methods, including the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

CLI Security Options

You can access the CLI using the console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console, Telnet, and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
  - Using RADIUS.See the"Configuring RADIUS" section on page 33-8.
  - Using TACACS+. See the"Configuring TACACS+" section on page 33-17.
- Local security control. See the "Local AAA Services" section on page 33-35.

These security features can also be configured for the following scenarios:

- iSCSI authentication (see the).
- Fibre Channel Security Protocol (FC-SP) authentication (see Chapter 37, "Configuring FC-SP and DHCHAP")

SNMP Security Options

The SNMP agent supports security features for SNMPv1, SNMPv2c, and SNMPv3. Normal SNMP security features apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

SNMP security options also apply to Fabric Manager and Device Manager.

See Chapter 32, "Configuring SNMP".

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for information on Fabric Manager and Device Manager.

Switch AAA Functionalities

Using the CLI or an SNMP application, you can configure AAA switch functionalities on any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

Authentication

Authentication is the process of verifying the identity of the person or device accessing the switch. This identity verification is based on the user ID and password combination provided by the entity trying to access the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

Note When you log in to a Cisco MDS switch successfully using the Fabric Manager or Device Manager through Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The switch authenticates the SNMPv3 protocol data units (PDUs) with your Telnet or SSH login name as the SNMPv3 user. The management station can temporarily use the Telnet or SSH login name as the SNMPv3 **auth** and **priv** passphrase. This temporary SNMP login is only allowed if you have one or more active MDS shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMPv3 operations.

Authorization

The following authorization roles exist in all Cisco MDS switches:

- Network operator (network-operator)—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (network-admin)— Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.
- Default-role—Has permission to use the GUI (Fabric Manager and Device Manager). This access is automatically granted to all users for accessing the GUI.

These roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.

Note If a user belongs only to one of the newly created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

Accounting

The accounting feature tracks and maintains a log of every management configuration used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over local AAA services:

- User password lists for each switch in the fabric can be managed more easily.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- The accounting log for all switches in the fabric can be centrally managed.
- User role mapping for each switch in the fabric can be managed more easily.

Remote Authentication Guidelines

If you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see Chapter 44, "Configuring IP Storage"). We recommend this method.
- SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for failover servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fail to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

AAA Service Configuration Options

AAA configuration in Cisco MDS 9000 Family switches is service based. You can have separate AAA configurations for the following services:

- Telnet or SSH login (Fabric Manager and Device Manager login)
- Console login
- iSCSI authentication (see)
- FC-SP authentication (see Chapter 37, "Configuring FC-SP and DHCHAP")
- Accounting

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the options fail, local is tried.

Caution Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local username with all numerics cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.

Note Even if local is not specified as one of the options, it is tried when all other configured options fail.

Table 33-1 provides the related CLI command for each AAA service configuration option.

Table 33-1 AAA Service Configuration Commands

AAA Service Configuration Option	Related Command
Telnet or SSH login (Cisco Fabric Manager and Device Manager login)	**aaa authentication login default**
Console login	**aaa authentication login console**
iSCSI authentication	**aaa authentication iscsi default**
FC-SP authentication	**aaa authentication dhchap default**
Accounting	**aaa accounting default**

Error-Enabled Status

When you log in, the login is processed by rolling over to local user database if the remote AAA servers do not respond. In this situation, the following message is displayed on the your screen—if you have enabled the error-enabled feature:

```
Remote AAA servers unreachable; local authentication done.
```

To enable this message display, use the **aaa authentication login error-enable** command.

To disable this message display, use the **no aaa authentication login error-enable** command.

To view the current display status, use the **show aaa authentication login error-enable** command (see Example 33-1).

Example 33-1 Displays AAA Authentication Login Information

```
switch# show aaa authentication login error-enable
enabled
```

AAA Server Monitoring

An unresponsive AAA server introduces a delay in the processing of AAA requests. An MDS switch can periodically monitor an AAA server to check whether it is responding (or alive) to save time in processing AAA requests. The MDS switch marks unresponsive AAA servers as dead and does not send AAA requests to any dead AAA servers. An MDS switch periodically monitors dead AAA servers and brings them to the alive state once they are responding. This monitoring process verifies that an AAA

server is in a working state before real AAA requests are sent its way. Whenever an AAA server changes to the dead or alive state, an SNMP trap is generated and the MDS switch warns the administrator that a failure is taking place before it can impact performance. See Figure 33-1 for AAA server states.

Figure 33-1 AAA Server States

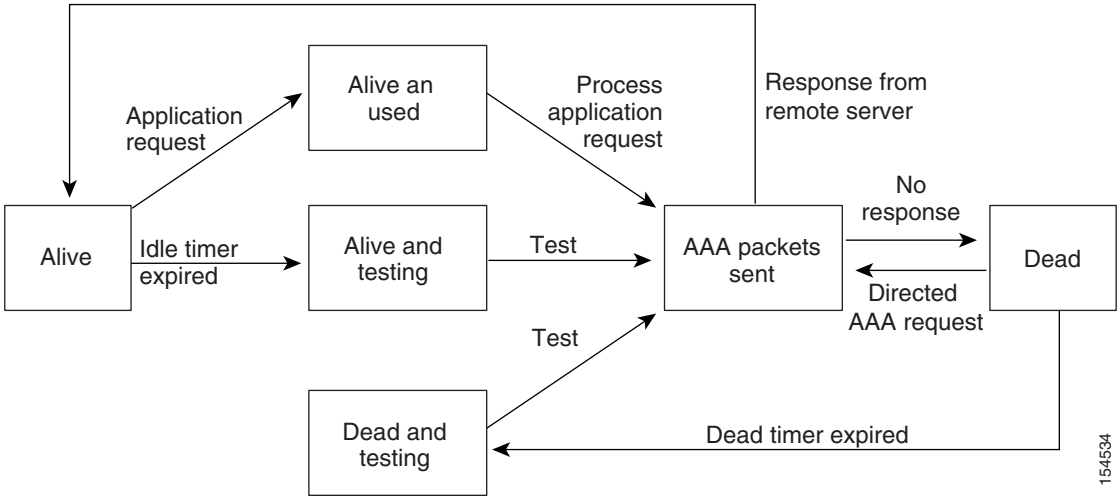

Note The monitoring interval for alive servers and dead servers is different and can be configured by the user. The AAA server monitoring is performed by sending a test authentication request to the AAA server.

The user name and password to be used in the test packet can be configured.

See the "Configuring RADIUS Server Monitoring Parameters" section on page 33-12 and"Displaying RADIUS Server Details" section on page 33-16.

Authentication and Authorization Process

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person managing the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

The following steps explain the authorization and authentication process:

Step 1 You can log in to the required switch in the Cisco MDS 9000 Family, using the Telnet, SSH, Fabric Manager/Device Manager, or console login options.

Step 2 When you have configured server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.

- If the AAA server fails to respond, then the next AAA server is contacted and so on until the remote server responds to the authentication request.
- If all AAA servers in the server group fail to respond, then the servers in the next server group are contacted.
- If all configured methods fail, then the local database is used for authentication.

Step 3 When you are successfully authenticated through a remote AAA server, then the following possible actions are taken:

- If the AAA server protocol is RADIUS, then user roles specified in the **cisco-av-pair** attribute are downloaded with an authentication response.
- If the AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.
- If user roles are not successfully retrieved from the remote AAA server, then the user is assigned the network-operator role.

Step 4 When your user name and password are successfully authenticated locally, you are allowed to log in, and you are assigned the roles configured in the local database.

Figure 33-2 shows a flow chart of the authorization and authentication process.

Figure 33-2 Switch Authorization and Authentication Flow

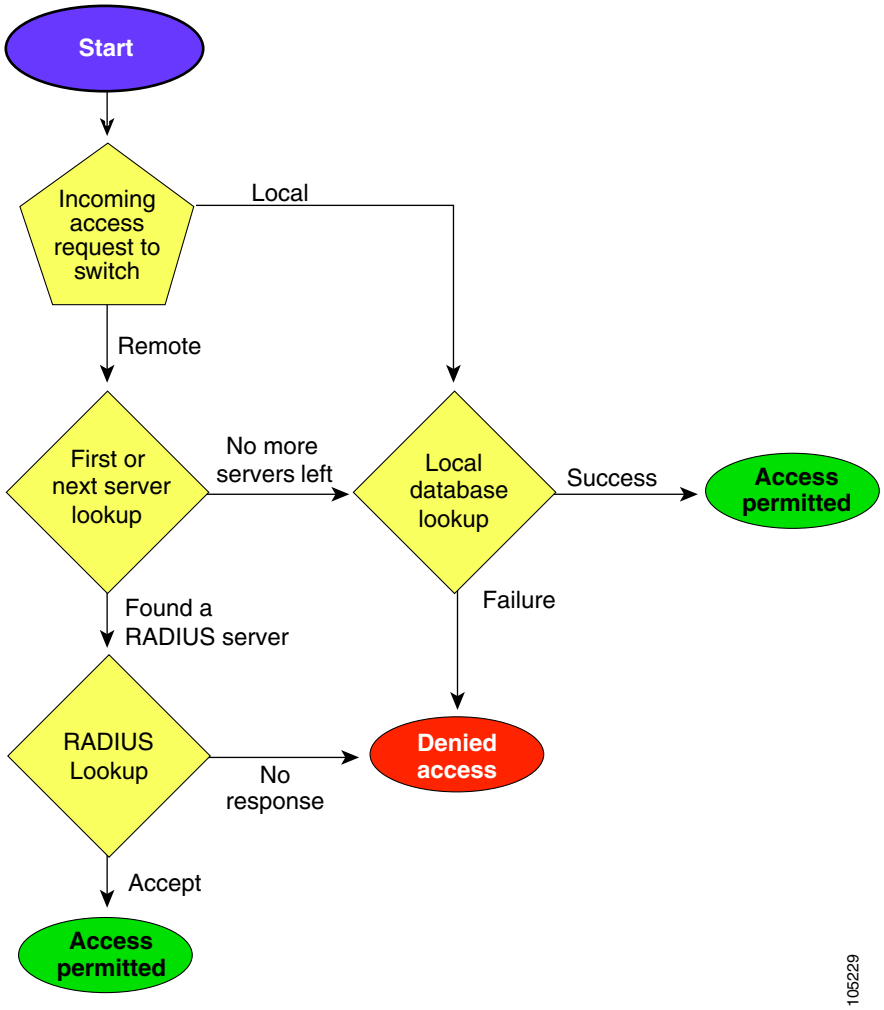

Note No more server groups left = no response from any server in all server groups.
No more servers left = no response from any server within this server group.

Configuring RADIUS

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

Setting the RADIUS Server Address

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys.

To specify the host RADIUS server IPv4 address and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server host 10.10.0.0 key HostKey`	Specifies the preshared key for the selected RADIUS server. This key overrides the key assigned using the **radius-server key** command. In this example, the host is 10.10.0.0 and the key is HostKey.
Step 3	`switch(config)# radius-server host 10.10.0.0 auth-port 2003`	Specifies the destination UDP port number to which the RADIUS authentication messages should be sent. In this example, the host is 10.10.0.0 and the authentication port is 2003. The default authentication port is 1812, and the valid range is 0 to 65366.
Step 4	`switch(config)# radius-server host 10.10.0.0 acct-port 2004`	Specifies the destination UDP port number to which RADIUS accounting messages should be sent. The default accounting port is 1813, and the valid range is 0 to 65366.
Step 5	`switch(config)# radius-server host 10.10.0.0 accounting`	Specifies this server to be used only for accounting purposes. **Note** If neither the **authentication** nor the **accounting** options are specified, the server is used for both accounting and authentication purposes.

	Command	Purpose
Step 6	`switch(config)# radius-server host 10.10.0.0 key 0 abcd`	Specifies a clear text key for the specified server. The key is restricted to 64 characters.
	`switch(config)# radius-server host 10.10.0.0 key 4 da3Asda2ioyuoiuH`	Specifies an encrypted key for the specified server. The key is restricted to 64 characters.

To specify the host RADIUS server IPv6 address and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server host 2001:0DB8:800:200C::417A Key HostKey`	Specifies the preshared key for the selected RADIUS server. This key overrides the key assigned using the **radius-server key** command. In this example, the host is 2001:0DB8:800:200C::417A and the key is HostKey.
Step 3	`switch(config)# radius-server host 2001:0DB8:800:200C::417A auth-port 2003`	Specifies the destination UDP port number to which the RADIUS authentication messages should be sent. In this example, the host is 2001:0DB8:800:200C::417A and the authentication port is 2003. The default authentication port is 1812, and the valid range is 0 to 65366.
Step 4	`switch(config)# radius-server host 2001:0DB8:800:200C::417A acct-port 2004`	Specifies the destination UDP port number to which RADIUS accounting messages should be sent. The default accounting port is 1813, and the valid range is 0 to 65366.
Step 5	`switch(config)# radius-server host 2001:0DB8:800:200C::417A accounting`	Specifies this server to be used only for accounting purposes. **Note** If neither the **authentication** nor the **accounting** options are specified, the server is used for both accounting and authentication purposes.
Step 6	`switch(config)# radius-server host 2001:0DB8:800:200C::417A key 0 abcd`	Specifies a clear text key for the specified server. The key is restricted to 64 characters.
	`switch(config)# radius-server host 2001:0DB8:800:200C::417A key 4 da3Asda2ioyuoiuH`	Specifies an encrypted key for the specified server. The key is restricted to 64 characters.

To specify the host RADIUS server DNS name and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server host radius2 key HostKey`	Specifies the preshared key for the selected RADIUS server. This key overrides the key assigned using the **radius-server key** command. In this example, the host is radius2 and the key is HostKey.

	Command	Purpose
Step 3	`switch(config)# r`**`adius-server host radius2 auth-port 2003`**	Specifies the destination UDP port number to which the RADIUS authentication messages should be sent. In this example, the host is radius2 and the authentication port is 2003. The default authentication port is 1812, and the valid range is 0 to 65366.
Step 4	`switch(config)#` **`radius-server host radius2 acct-port 2004`**	Specifies the destination UDP port number to which RADIUS accounting messages should be sent. The default accounting port is 1813, and the valid range is 0 to 65366.
Step 5	`switch(config)#` **`radius-server host radius2 accounting`**	Specifies this server to be used only for accounting purposes. **Note** If neither the **authentication** nor the **accounting** options are specified, the server is used for both accounting and authentication purposes.
Step 6	`switch(config)#` **`radius-server host radius2 key 0 abcd`**	Specifies a clear text key for the specified server. The key is restricted to 64 characters.
	`switch(config)#` **`radius-server host radius2 key 4 da3Asda2ioyuoiuH`**	Specifies an encrypted key for the specified server. The key is restricted to 64 characters.

About the Default RADIUS Server Encryption Type and Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option in the **radius-server host** command.

Configuring the Default RADIUS Server Encryption Type and Preshared Key

To configure the RADIUS preshared key, follow these steps:

	Command	Purpose
Step 1	`switch#` **`config t`**	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# radius-server key AnyWord`	Configures a preshared key (AnyWord) to authenticate communication between the RADIUS client and server. The default is clear text.
	`switch(config)# radius-server key 0 AnyWord`	Configures a preshared key (AnyWord) specified in clear text (indicated by 0) to authenticate communication between the RADIUS client and server.
	`switch(config)# radius-server key 7 abe4DFeeweo00o`	Configures a preshared key (specified in encrypted text) specified in encrypted text (indicated by 7) to authenticate communication between the RADIUS client and server.

Setting the RADIUS Server Timeout Interval

You can configure a global timeout value between transmissions for all RADIUS servers.

Note If timeout values are configured for individual servers, those values override the globally configured values.

To specify the timeout values between retransmissions to the RADIUS servers, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server timeout 30`	Configures the global timeout period in seconds for the switch to wait for a response from all TACACS+ servers before the switch declares a timeout failure. The time ranges from 1 to 1440 seconds.
	`switch(config)# no radius-server timeout 30`	Reverts the transmission time to the default value (1 second).

Setting Transmission Retry Count for the RADIUS Server

By default, a switch retries transmission to a RADIUS server only once before reverting to local authentication. You can increase this number up to a maximum of five retries per server.To specify the number of times that RADIUS servers should try to authenticate a user, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server retransmit 3`	Configures the number of times (3) the switch tries to connect to a RADIUS server(s) before reverting to local authentication.
	`switch(config)# no radius-server retransmit`	Reverts to the default retry count (1).

Configuring RADIUS Server Monitoring Parameters

You can configure parameters for monitoring RADIUS servers. You can configure this option to test the server periodically, or you can run a one-time only test.

This section includes the following topics:

Configuring the Test Idle Timer

The test idle timer specifies the interval during which a RADIUS server receives no requests before the MDS switch sends out a test packet.

Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic RADIUS server monitoring is not performed.

To configure the idle timer, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server host 10.1.1.1 test idle-time 20`	Configures the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.
Step 3	`switch(config)# no radius-server host 10.1.1.1 test idle-time 20`	Reverts to the default value (0 minutes).

Configuring Test User Name

You can configure a username and password for periodic RADIUS server status testing. You do not need to configure the test username and password to issue test messages to monitor RADIUS servers. You can use the default test username (test) and default password (test).

Note We recommend that the test username not be the same as an existing username in the RADIUS database for security reasons.

To configure the optional username and password for periodic RADIUS server status testing, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# radius-server host 10.1.1.1 test username testuser`	Configures the test user (testuser) with the default password (test). The default user name is test.
	`switch(config)# no radius-server host 10.1.1.1 test username testuser`	Removes the test user name (testuser).
	`switch(config)# radius-server host 10.1.1.1 test username testuser password Ur2Gd2BH`	Configures the test user (testuser) and assigns a strong password. For guidelines for creating strong passwords, see the "Characteristics of Strong Passwords" section on page 31-12.

Configuring the Dead Timer

The dead timer specifies the interval that the MDS switch waits, after declaring that a RADIUS server is dead, before sending out a test packet to determine if the server is now alive.

Note The default dead timer value is 0 minutes. When the dead timer interval is 0 minutes, RADIUS server monitoring is not performed unless the RADIUS server is part of a server group and the dead-time interval for the group is greater than 0 minutes. (See the "Server Groups" section on page 33-4.)

Note If the dead timer of a dead RADIUS server expires before it is sent a RADIUS test message, that server is marked as alive again even if it is still not responding. To avoid this scenario, configure a test user with a shorter idle time than the dead timer time.

To configure the dead timer, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server deadtime 30`	Configures the dead timer interval value in minutes. The valid range is 1 to 1440 minutes.
Step 3	`switch(config)# no radius-server deadtime 30`	Reverts to the default value (0 minutes).

Sending RADIUS Test Messages for Monitoring

You can manually send test messages to monitor a RADIUS server.

To send the test message to the RADIUS server, follow this step:

	Command	Purpose
Step 1	`switch# test aaa server radius 10.10.1.1 test test`	Sends a test message to a RADIUS server using the default username (test) and password (test).
	`switch# test aaa server radius 10.10.1.1 testuser Ur2Gd2BH`	Sends a test message to a RADIUS server using a configured test username (testuser) and password (Ur2Gd2BH). **Note** A configured username and password is optional (see the "Configuring Test User Name" section on page 33-12).

About Users Specifying a RADIUS Server at Login

By default, an MDS switch forwards an authentication request to the first server in the RADIUS server group. You can configure the switch to allow the user to specify which RADIUS server to send the authenticate request by enabling the directed request option. If you enable this option, the user can log in as *username@hostname*, where the *hostname* is the name of a configured RADIUS server.

Allowing Users to Specify a RADIUS Server at Login

To allow users logging into an MDS switch to select a RADIUS server for authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius-server directed-request`	Allows users to specify a RADIUS server to send the authentication request when logging in.
	`switch(config)# no radius-server directed-request`	Reverts to sending the authentication request to the first server in the server group (default).

You can use the **show tacacs-server directed-request** command to display the RADIUS directed request configuration.

```
switch# show radius-server directed-request
disabled
```

About Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for

general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named **cisco-avpair.** The value is a string with the following format:

```
protocol : attribute seperator value *
```

Where **protocol** is a Cisco attribute for a particular type of authorization, **separator** is `=` (equal sign) for mandatory attributes, and `*` (asterisk) is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco SAN-OS software:

- **Shell** protocol—used in Access-Accept packets to provide user profile information.
- **Accounting** protocol—used in Accounting-Request packets. If a value contains any white spaces, it should be put within double quotation marks.

The following attributes are supported by the Cisco SAN-OS software:

- **roles**—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles **vsan-admin** and **storage-admin**, the value field would be "**vsan-admin storage-admin"**. This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

  ```
  shell:roles="network-admin vsan-admin"
  shell:roles*"network-admin vsan-admin"
  ```

  When an VSA is specified as **shell:roles*"network-admin vsan-admin"**, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- **accountinginfo**—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying SNMPv3 on AAA Servers

The vendor/custom attribute **cisco-av-pair** can be used to specify user's role mapping using the format:

```
shell:roles="roleA roleB …"
```

If the roll option in the **cisco-av-pair** attribute is not set, the default user role is network-operator.

The VSA format optionally specifies your SNMPv3 authentication and privacy protocol attributes also as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the **cisco-av-pair** attribute on the ACS server, MD5 and DES are used by default.

Displaying RADIUS Server Details

Use the **show radius-server** command to display configured RADIUS parameters as shown in Example 33-2.

Example 33-2 Displays Configured RADIUS Information

```
switch# show radius-server
Global RADIUS shared secret:*******
retransmission count:5
timeout value:10
following RADIUS servers are configured:
        myradius.cisco.users.com:
                available for authentication on port:1812
                available for accounting on port:1813
        172.22.91.37:
                available for authentication on port:1812
                available for accounting on port:1813
                RADIUS shared secret:******
        10.10.0.0:
                available for authentication on port:1812
                available for accounting on port:1813
                RADIUS shared secret:******
```

Example 33-3 Displays Configured RADIUS Server-Group Order

```
switch# show radius-server groups
total number of groups:4
following RADIUS server groups are configured:
        group radius:
                server: all configured radius servers
        group Group1:
                server: Server3 on auth-port 1812, acct-port 1813
                server: Server5 on auth-port 1812, acct-port 1813
        group Group5:
```

Displaying RADIUS Server Statistics

You can display RADIUS server statistics using the **show radius-server statistics** command.

Example 33-4 Displays RADIUS Server Statistics

```
switch# show radius-server statistics 10.1.3.2
Server is not monitored

Authentication Statistics
        failed transactions: 0
        sucessfull transactions: 0
        requests sent: 0
        requests timed out: 0
        responses with no matching requests: 0
        responses not processed: 0
        responses containing errors: 0
```

```
Accounting Statistics
        failed transactions: 0
        sucessfull transactions: 0
        requests sent: 0
        requests timed out: 0
        responses with no matching requests: 0
        responses not processed: 0
        responses containing errors: 0
```

Configuring TACACS+

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

This section includes the following topics:

About TACACS+

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The TACACS+ has the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- Uses the TCP transport protocol to send data between the AAA client and server, making reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

About TACACS+ Server Default Configuration

Fabric Manager allows you to set up a default configuration that can be used for any TACACS+ server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Preshared key
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a TACACS+ server at login

About the Default TACACS+ Server Encryption Type and Preshared Key

You need to configure the TACACS+ preshared key to authenticate the switch to the TACACS+ server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all TACACS+ server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring and individual TACACS+ server.

Enabling TACACS+

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the TACACS+ feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

To enable TACACS+ for a Cisco MDS switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs+ enable`	Enables the TACACS+ in this switch.
	`switch(config)# no tacacs+ enable`	Disables (default) the TACACS+ in this switch.

Setting the TACACS+ Server Address

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server (see the "Setting the Timeout Value" section on page 33-21).

Note You can use the dollar sign ($) and the percent sign (%) in global secret keys.

To configure the TACACS+ server IPv4 address and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server host 171.71.58.91`	Configures the TACACS+ server identified by the specified IPv4 address.
	`switch(config)# no tacacs-server host 171.71.58.91`	Deletes the specified TACACS+ server identified by the IPv4 address. By default, no server is configured.
Step 3	`switch(config)# tacacs-server host 171.71.58.91 port 2`	Configures the TCP port for all TACACS+ requests.
	`switch(config)# no tacacs-server host 171.71.58.91 port 2`	Reverts to the factory default of using port 49 for server access.
Step 4	`switch(config)# tacacs-server host 171.71.58.91 key MyKey`	Configures the TACACS+ server identified by the specified domain name and assigns the secret key.
Step 5	`switch(config)# tacacs-server host 171.71.58.91 timeout 25`	Configures the timeout period for the switch to wait for a response from the specified server before it declares a timeout failure.

To configure the TACACS+ server IPv6 address and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server host 2001:0DB8:800:200C::417A` `warning: no key is configured for the host`	Configures the TACACS+ server identified by the specified IPv6 address.
	`switch(config)# no tacacs-server host 2001:0DB8:800:200C::417A`	Deletes the specified TACACS+ server identified by the IPv6 address. By default, no server is configured.
Step 3	`switch(config)# tacacs-server host 2001:0DB8:800:200C::417A port 2`	Configures the TCP port for all TACACS+ requests.
	`switch(config)# no tacacs-server host 2001:0DB8:800:200C::417A port 2`	Reverts to the factory default of using port 49 for server access.
Step 4	`switch(config)# tacacs-server host 2001:0DB8:800:200C::417A key MyKey`	Configures the TACACS+ server identified by the specified domain name and assigns the secret key.
Step 5	`switch(config)# tacacs-server host 2001:0DB8:800:200C::417A timeout 25`	Configures the timeout period for the switch to wait for a response from the specified server before it declares a timeout failure.

To configure the TACACS+ server DNS name and other options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server host host1.cisco.com` `warning: no key is configured for the host`	Configures the TACACS+ server identified by the specified DNS name.
	`switch(config)# no tacacs-server host host1.cisco.com`	Deletes the specified TACACS+ server identified by the DNS name. By default, no server is configured.

	Command	Purpose
Step 3	`switch(config)# tacacs-server host host1.cisco.com port 2`	Configures the TCP port for all TACACS+ requests.
	`switch(config)# no tacacs-server host host1.cisco.com port 2`	Reverts to the factory default of using port 49 for server access.
Step 4	`switch(config)# tacacs-server host host1.cisco.com key MyKey`	Configures the TACACS+ server identified by the specified domain name and assigns the secret key.
Step 5	`switch(config)# tacacs-server host host1.cisco.com timeout 25`	Configures the timeout period for the switch to wait for a response from the specified server before it declares a timeout failure.

Setting the Global Secret Key

You can configure global values for the secret key for all TACACS+ servers.

Note If secret keys are configured for individual servers, those keys override the globally configured key.

Note You can use the dollar sign ($) and the percent sign (%) in global secret keys.

To set the secret key for TACACS+ servers, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server key 7 3sdaA3daKUngd`	Assigns the global secret key (in encrypted format) to access the TACACS+ server. This example specifies **7** to indicate the encrypted format being used. If this global key and the individual server keys are not configured, clear text messages are sent to the TACACS+ server(s).
	`switch(config)# no tacacs-server key oldPword`	Deletes the configured global secret key to access the TACACS+ server and reverts to the factory default of allowing access to all configured servers.

Setting the Timeout Value

You can configure a global timeout value between transmissions for all TACACS+ servers.

Note If timeout values are configured for individual servers, those values override the globally configured values.

To set the global timeout value for TACACS+ servers, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server timeout 30`	Configures the global timeout period in seconds for the switch to wait for a response from all TACACS+ servers before the switch declares a timeout failure. The time ranges from 1 to 1440 seconds.
	`switch(config)# no tacacs-server timeout 30`	Deletes the configured timeout period and reverts to the factory default of 5 seconds.

About TACACS+ Servers

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. Fabric Manager or Device Manager enables the TACACS+ feature automatically when you configure a TACACS+ server.

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server.

Note Prior to Cisco MDS SAN-OS Release 2.1(2), you can use the dollar sign ($) in the key but the key must be enclosed in double quotes, for example “k$”. The percent sign (%) is not allowed. In Cisco MDS SAN-OS Release 2.1(2) and later, you can use the dollar sign ($) without double quotes and the percent sign (%) in global secret keys.

You can configure global values for the secret key for all TACACS+ servers.

Note If secret keys are configured for individual servers, those keys override the globally configured key.

Configuring TACACS+ Server Monitoring Parameters

You can configure parameters for monitoring TACACS+ servers.

This section includes the following topics:

Configuring the TACACS+ Test Idle Timer

The test idle timer specifies the interval during which a TACACS+ server receives no requests before the MDS switch sends out a test packet.

Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic TACACS+ server monitoring is not performed.

To configure the idle timer, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server host 10.1.1.1 test idle-time 20`	Configures the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.
Step 3	`switch(config)# no tacacs-server host 10.1.1.1 test idle-time 20`	Reverts to the default value (0 minutes).

Configuring Test Username

You can configure a username and password for periodic TACACS+ server status testing. You do not servers. You can use the default test username (test) and default password (test).

To configure the optional username and password for periodic TACACS+ server status testing, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server host 10.1.1.1 test username testuser`	Configures the test user (testuser) with the default password (test). The default username is test.
	`switch(config)# no tacacs-server host 10.1.1.1 test username testuser`	Removes the test user (testuser).
	`switch(config)# tacacs-server host 10.1.1.1 test username testuser password Ur2Gd2BH`	Configures the test user (testuser) and assigns a strong password. For guidelines for creating strong passwords, see the "Characteristics of Strong Passwords" section on page 31-12.

Configuring the Dead Timer

The dead timer specifies the interval that the MDS switch waits, after declaring a TACACS+ server is dead, before sending out a test packet to determine if the server is now alive.

Note The default dead timer value is 0 minutes. TACACS+ server monitoring is not performed if the dead timer interval is 0 minutes, unless the TACACS+ server is a part of a bigger group with the dead-time interval greater than 0 minutes. (See "Configuring RADIUS" section on page 33-8.)

Note If the dead timer of a dead TACACS+ server expires before it is sent a TACACS+ test message, that server is marked as alive again even if it is still not responding. To avoid this scenario, configure a test user with a shorter idle time than the dead timer time.

To configure the dead timer, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server deadtime 30`	Configures the dead-time interval value in minutes. The valid range is 1 to 1440 minutes.
Step 3	`switch(config)# no tacacs-server deadtime 30`	Reverts to the default value (0 minutes). **Note** When the dead-time interval is 0 minutes, TACACS+ server monitoring is not performed unless the TACACS+ server is part of a server group and the dead-time interval for the group is greater than 0 minutes. (See the "Configuring RADIUS" section on page 33-8.)

Sending TACACS+ Test Messages for Monitoring

You can manually send test messages to monitor a TACACS+ server.

To send the test message to the TACACS+ server, follow these steps:

Command	Purpose
`switch# test aaa server tacacs+ 10.10.1.1 test test`	Sends a test message to a TACACS+ server using the default username (test) and password (test).
`switch# test aaa server tacacs+ 10.10.1.1 testuser Ur2Gd2BH`	Sends a test message to a TACACS+ server using a configured test username and password. A configured username and password is optional (see the "Configuring Test Username" section on page 33-22).

Password Aging Notification through TACACS+ Server

Password aging notification is initiated when the user authenticates to a Cisco MDS 9000 switch via a TACACS+ account. The user is notified when a password is about to expire or has expired. If the password has expired, user is prompted to change the password.

Note As of Cisco MDS SAN-OS Release 3.2(1), only TACACS+ supports password aging notification. If you try to use RADIUS servers by enabling this feature, RADIUSs will generate a SYSLOG message and authentication will fall back to the local database.

Password aging notification facilitates the following:

- Password change — You can change your password by entering a blank password.
- Password aging notification — Notifies password aging. Notification happens only if the AAA server is configured.
- Password change after expiration — Initiates password change after the old password expires. Initiation happens from the AAA server.

To enable the password aging option in the AAA server, enter the following command:

```
aaa authentication login password-aging enable
```

To determine whether or not password aging notification is enabled or disabled in the AAA server, enter the following command:

```
show aaa authentication login password-aging
```

About Users Specifying a TACACS+ Server at Login

By default, an MDS switch forwards an authentication request to the first server in the TACACS+ server group. You can configure the switch to allow the user to specify which TACACS+ server to send the authenticate request. If you enable this feature, the user can log in as *username@hostname*, where the *hostname* is the name of a configured TACACS+ server.

Allowing Users to Specify a TACACS+ Server at Login

To allow users logging into an MDS switch to select a TACACS+ server for authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs-server directed-request`	Allows users to specify a TACACS+ server to send the authentication request when logging in.
	`switch(config)# no tacacs-server directed-request`	Reverts to sending the authentication request to the first server in the server group (default).

You can use the **show tacacs-server directed-request** command to display the TACACS+ directed request configuration.

```
switch# show tacacs-server directed-request
disabled
```

Defining Custom Attributes for Roles

Cisco MDS 9000 Family switches use the TACACS+ custom attribute for service shells to configure roles to which a user belongs. TACACS+ attributes are specified in **name=value** format. The attribute name for this custom attribute is **cisco-av-pair**. The following example illustrates how to specify roles using this attribute:

```
cisco-av-pair=shell:roles="network-admin vsan-admin"
```

You can also configure optional custom attributes to avoid conflicts with non-MDS Cisco switches using the same AAA servers.

```
cisco-av-pair*shell:roles="network-admin vsan-admin"
```

Additional custom attribute shell:roles are also supported:

```
shell:roles="network-admin vsan-admin"
```

or

```
shell:roles*"network-admin vsan-admin"
```

Note TACACS+ custom attributes can be defined on an Access Control Server (ACS) for various services (for example, shell). Cisco MDS 9000 Family switches require the TACACS+ custom attribute for the service shell to be used for defining roles.

Supported TACACS+ Server Parameters

The Cisco SAN-OS software currently supports the following parameters for the listed TACACS+ servers:

- TACACS+

  ```
  cisco-av-pair=shell:roles="network-admin"
  ```

- Cisco ACS TACACS+

```
shell:roles="network-admin"
shell:roles*"network-admin"
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair*shell:roles*"network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

- Open TACACS+

```
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

Displaying TACACS+ Server Details

Use the **show aaa** and **show tacacs-server** commands to display information about TACACS+ server configuration in all switches in the Cisco MDS 9000 Family as shown in Examples 33-5 to 33-10.

Example 33-5 Displays Configured TACACS+ Server Information

```
switch# show tacacs-server
Global TACACS+ shared secret:***********
timeout value:30
total number of servers:3

following TACACS+ servers are configured:
        171.71.58.91:
                available on port:2
        cisco.com:
                available on port:49
        171.71.22.95:
                available on port:49
                TACACS+ shared secret:*****
```

Example 33-6 Displays AAA Authentication Information

```
switch# show aaa authentication
         default: group TacServer local none
         console: local
         iscsi: local
         dhchap: local
```

Example 33-7 Displays AAA Authentication Login Information

```
switch# show aaa authentication login error-enable
enabled
```

Example 33-8 Displays Configured TACACS+ Server Groups

```
switch# show tacacs-server groups
total number of groups:2

following TACACS+ server groups are configured:
        group TacServer:
                server 171.71.58.91 on port 2
        group TacacsServer1:
                server ServerA on port 49
                server ServerB on port 49:
```

Example 33-9 Displays All AAA Server Groups

```
switch# show aaa groups
radius
TacServer
```

Example 33-10 Displays TACACS+ Server Statistics

```
switch# show tacacs-server statistics 10.1.2.3
Server is not monitored

Authentication Statistics
        failed transactions: 0
        sucessfull transactions: 0
        requests sent: 0
        requests timed out: 0
        responses with no matching requests: 0
        responses not processed: 0
        responses containing errors: 0

Authorization Statistics
        failed transactions: 0
        sucessfull transactions: 0
        requests sent: 0
        requests timed out: 0
        responses with no matching requests: 0
        responses not processed: 0
        responses containing errors: 0

Accounting Statistics
        failed transactions: 0
        sucessfull transactions: 0
        requests sent: 0
        requests timed out: 0
        responses with no matching requests: 0
        responses not processed: 0
        responses containing errors: 0
```

Configuring Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol, either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

The AAA server monitoring feature can mark an AAA server as dead. You can configure a period of time in minutes to elapse before the switch sends requests to a dead AAA server. (See the "AAA Server Monitoring" section on page 33-5.)

You can configure these server groups at any time but they only take effect when you apply them to an AAA service. You configure AAA policies for CLI users or Fabric Manager or Device Manager users.

To configure a RADIUS server group, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# aaa group server radius RadServer` `switch(config-radius)#`	Creates a server group named RadServer and enters the RADIUS server group configuration submode for that group.
	`switch(config)# no aaa group server radius RadServer`	Deletes the server group called RadServer from the authentication list.
Step 3	`switch(config-radius)# server 10.71.58.91`	Configures the RADIUS server at IPv4 address 10.71.58.91 to be tried first within the server group RadServer. **Tip** If the specified RADIUS server is not found, configure it using the **radius-server host** command and retry this command.
Step 4	`switch(config-radius)# server 2001:0DB8:800:200C::417A`	Configures the RADIUS server at IPv6 address 2001:0DB8:800:200C::417A to be tried first within the server group RadServer.
	`switch(config-radius)# no server 2001:0DB8:800:200C::417A`	Removes the RADIUS server at IPv6 address 2001:0DB8:800:200C::417A from the server group RadServer.
Step 5	`switch(config-radius)# exit`	Returns to configuration mode.
Step 6	`switch(config)# aaa group server radius RadiusServer` `switch(config-radius)#`	Creates a server group named RadiusServer and enters the RADIUS server group configuration submode for that group.
Step 7	`switch(config-radius)# server ServerA`	Configures ServerA to be tried first within the server group called the RadiusServer1. **Tip** If the specified RADIUS server is not found, configure it using the **radius-server host** command and retry this command.

	Command	Purpose
Step 8	`switch(config-radius)# server ServerB`	Configures ServerB to be tried second within the server group RadiusServer1.
Step 9	`switch(config-radius)# deadtime 30`	Configures the monitoring dead time to 30 minutes. The range is 0 through 1440. **Note** If the dead-time interval for an individual RADIUS server is greater than 0, that value takes precedence over the value set for the server group.
	`switch(config-radius)# no deadtime 30`	Reverts to the default value (0 minutes). **Note** If the dead-time interval for both the RADIUS server group and an individual TACACS+ server in the RADIUS server group is set to 0, the switch does not mark the RADIUS server as dead when it is found to be unresponsive by periodic monitoring. Also, the switch does not perform dead server monitoring for that RADIUS server. (See the "Configuring RADIUS Server Monitoring Parameters" section on page 33-12.)

To verify the configured server group order, use the **show radius-server groups** command:

```
switch# show radius-server groups
total number of groups:2

following RAIDUS server groups are configured:
        group RadServer:
                server 10.71.58.91 on port 2
        group RadiusServer1:
                server ServerA on port 49
                server ServerB on port 49:
```

To configure a TACACS+ server group, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# aaa group server tacacs+ TacacsServer1` `switch(config-tacacs+)#`	Creates a server group named TacacsServer1 and enters the submode for that group.
	`switch(config)# no aaa group server tacacs+ TacacsServer1`	Deletes the server group called TacacsServer1 from the authentication list.
Step 3	`switch(config-tacacs+)# server ServerA`	Configures ServerA to be tried first within the server group called the TacacsServer1. **Tip** If the specified TACACS+ server is not found, configure it using the **tacacs-server host** command and retry this command.

	Command	Purpose
Step 4	`switch(config-tacacs+)# server ServerB`	Configures ServerB to be tried second within the server group TacacsServer1.
	`switch(config-tacacs+)# no server ServerB`	Deletes ServerB within the TacacsServer1 list of servers.
Step 5	`switch(config-tacacs+)# deadtime 30`	Configures the monitoring dead time to 30 minutes. The range is 0 through 1440. **Note** If the dead-time interval for an individual TACACS+ server is greater than 0, that value takes precedence over the value set for the server group.
	`switch(config-tacacs+)# no deadtime 30`	Reverts to the default value (0 minutes). **Note** If the dead-time interval for both the TACACS+ server group and an individual TACACS+ server in the TACACS+ server group is set to 0, the switch does not mark the TACACS+ server as dead when it is found to be unresponsive by periodic monitoring. Also, the switch does not perform dead server monitoring for that TACACS+ server. (See the "Configuring TACACS+ Server Monitoring Parameters" section on page 33-21.)

AAA Server Distribution

Configuration for RADIUS and TACACS+ AAA on an MDS switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default (see Chapter 6, "Using the CFS Infrastructure").

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.

Note Server group configurations are not distributed.

Note For an MDS switch to participate in AAA server configuration distribution, it must be running Cisco MDS SAN-OS Release 2.0(1b) or later.

Enabling AAA Server Distribution

Only switches where distribution is enabled can participate in the distribution activity.

To enable RADIUS server distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius distribute`	Enables RADIUS configuration distribution in this switch.
	`switch(config)# no radius distribute`	Disables RADIUS configuration distribution in this switch (default).

To enable TACACS+ server distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs+ distribute`	Enables TACACS+ configuration distribution in this switch.
	`switch(config)# no tacacs+ distribute`	Disables TACACS+ configuration distribution in this switch (default).

Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a RADIUS/TACACS+ server or global configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.

Note After you issue the first configuration command related to AAA servers, all server and global configurations that are created (including the configuration that caused the distribution session start) are stored in a temporary buffer, not in the running configuration.

Displaying the Session Status

Once the implicit distribution session has started, you can check the session status **.** You see the **distribution status** on the CFS tabuse the **show radius** command.

```
switch# show radius distribution status
distribution : enabled
session ongoing: yes
session owner: admin
session db: exists
merge protocol status: merge activation done


last operation: enable
last operation status: success
```

Once the implicit distribution session has started, you can check the session status using the **show tacacs+ distribution status** command.

```
switch# show tacacs+ distribution status
distribution : enabled
session ongoing: yes
session owner: admin
session db: exists
merge protocol status: merge activation done


last operation: enable
last operation status: success
```

Displaying the Pending Configuration

To display the RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer use the **show radius pending** command, follow these steps:

```
switch(config)# show radius pending-diff
   +radius-server host testhost1 authentication accounting
   +radius-server host testhost2 authentication accounting
```

To display the TACACS+ global and/or server configuration stored in the temporary buffer, use the **show tacacs+ pending** command.

```
switch(config)# show tacacs+ pending-diff
   +tacacs-server host testhost3
   +tacacs-server host testhost4
```

Committing the Distribution

The RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer can be applied to the running configuration across all switches in the fabric (including the originating switch).

To commit RADIUS configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius commit`	Commits the RADIUS configuration changes to the running configuration.

To commit TACACS+ configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs+ commit`	Commits the TACACS+ configuration changes to the running configuration.

Discarding the Distribution Session

Discarding the distribution of a session in progress causes the configuration in the temporary buffer to be dropped. The distribution is not applied.

To discard the RADIUS session-in-progress distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# radius abort`	Discards the RADIUS configuration changes to the running configuration.

To discard the TACACS+ session-in-progress distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tacacs+ abort`	Discards the TACACS+ configuration changes to the running configuration.

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the RADIUS feature, enter the **clear radius session** command from any switch in the fabric.

```
switch# clear radius session
```

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the TACACS+ feature, enterthe **clear tacacs+ session** command from any switch in the fabric.

```
switch# clear tacacs+ session
```

Merge Guidelines for RADIUS and TACACS+ Configurations

The RADIUS and TACACS+ server and global configuration are merged when two fabrics merge. The merged configuration is applied to CFS distribution-enabled switches.

When merging the fabric, be aware of the following conditions:

- The server groups are not merged.
- The server and global keys are not changed during the merge.
- The merged configuration contains all servers found on all CFS enabled switches.
- The timeout and retransmit parameters of the merged configuration are the largest values found per server and global configuration.

Caution If there is a conflict between two switches in the server ports configured, the merge fails.

Use the **show radius distribution status** command to view the status of the RADIUS fabric merge as shown in Example 33-11.

Example 33-11 Displays the RADIUS Fabric Merge Status

```
switch# show radius distribution status
distribution : enabled
session ongoing: no
session db: does not exist
merge protocol status: merge response received
merge error: conflict: server dmtest2 has auth-port 1812 on this switch and 1999
on remote

last operation: enable
last operation status: success
```

Use the **show tacacs+ distribution status** command to view the status of the TACACS+ fabric merge as shown in Example 33-12.

Example 33-12 Displays the TACACS+ Fabric Merge Status

```
switch# show tacacs+ distribution status
distribution : enabled
session ongoing: no
session db: does not exist
merge protocol status: merge activation done

last operation: enable
last operation status: success
```

MSCHAP Authentication

Microsoft Challenge Handshake Authentication Protocol (MSCHAP) is the Microsoft version of CHAP. You can use MSCHAP for user logins to an MDS switch through a remote authentication server (RADIUS or TACACS+).

About Enabling MSCHAP

By default, the switch uses Password Authentication Protocol (PAP) authentication between the switch and the remote server. If you enable MSCHAP, you need to configure your RADIUS server to recognize the MSCHAP vendor-specific attributes. See the "About Vendor-Specific Attributes" section on page 33-14. Table 39-1 shows the RADIUS vendor-specific attributes required for MSCHAP.

To enable MSCHAP authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# aaa authentication login mschap enable`	Enables MSCHAP login authentication.

You can use the **show aaa authentication login mschap** command to display the MSCHAP authentication configuration.

```
switch# show aaa authentication login mschap
disabled
```

Local AAA Services

The system maintains the username and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information. Use the **username** command to configure local users and their roles (see the "Configuring User Accounts" section on page 31-11).

Use the **show accounting log** command to view the local accounting log as shown in Example 33-13.

Example 33-13 Displays the Accounting Log Information

```
switch# show accounting log

Sat Jan 24 03:22:06 1981:stop:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:start:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:update:snmp_349154526_171.71.58.69:admin:Added member [
 WWN: 21:00:00:20:37:a6:be:00 ID: 2] to zone test-27 on VSAN 1
...
Sat Jan 24 23:59:56 1981:stop:/dev/pts/0_349228792:root:shell terminated
Sun Jan 25 00:00:06 1981:start:/dev/pts/1_349228806:admin:
```

Disabling AAA Authentication

You can turn off password verification using the **none** option. If you configure this option, users can log in without giving a valid password. But the user should at least exist locally on the Cisco MDS 9000 Family switch.

Caution Use this option cautiously. If configured, any user can access the switch at any time.

Use the **none** option in the **aaa authentication login** command to disable password verification.

A user created by entering the **username** command will exist locally on the Cisco MDS 9000 Family switch.

Displaying AAA Authentication

The **show aaa authentication** command displays the configured authentication methods as shown in Example 33-14.

Example 33-14 Displays Authentication Information

```
switch# show aaa authentication

    No AAA Authentication
    default: group TacServer local none
    console: local none
    iscsi: local
    dhchap: local
```

Configuring Accounting Services

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting and auditing purposes. Accounting can be implemented locally or remotely (using RADIUS). The default maximum size of the accounting log is 250,000 bytes and cannot be changed.

Tip The Cisco MDS 9000 Family switch uses interim-update RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have log update/watchdog packets flags in the AAA client configuration. Turn on this flag to ensure proper RADIUS accounting.

Note Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

Displaying Accounting Configuration

To display configured accounting information use **show accounting** command. See Examples 33-15 to 33-17. To specify the size of the local accounting log to be displayed, use the **show accounting log** command. By default about 250 KB of accounting log is displayed.

Example 33-15 Displays Two Samples of Configured Accounting Parameters

```
switch# show accounting config
show aaa accounting
                default: local

switch# show aaa accounting
                default: group rad1
```

Example 33-16 Displays 60,000 Bytes of the Accounting Log

```
switch# show accounting log 60000
Fri Jan 16 15:28:21 1981:stop:snmp_348506901_64.104.131.208:admin:
Fri Jan 16 21:17:04 1981:start:/dev/pts/0_348527824:admin:
Fri Jan 16 21:35:45 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 21:58:17 1981:start:snmp_348530297_171.71.150.105:admin:
...
```

Example 33-17 Displays the Entire Log File

```
switch# show accounting log
Fri Jan 16 15:28:21 1981:stop:snmp_348506901_64.104.131.208:admin:
Fri Jan 16 21:17:04 1981:start:/dev/pts/0_348527824:admin:
Fri Jan 16 21:35:45 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 21:58:17 1981:start:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:17 1981:stop:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:start:snmp_348530298_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:stop:snmp_348530298_171.71.150.105:admin:
...
Fri Jan 16 23:37:02 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 23:37:26 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Fri Jan 16 23:53:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server3
Fri Jan 16 23:54:00 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server5
Fri Jan 16 23:54:22 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerA
Fri Jan 16 23:54:25 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerB
Fri Jan 16 23:55:03 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Sat Jan 17 00:01:41 1981:start:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:41 1981:stop:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:start:snmp_348537702_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:stop:snmp_348537702_171.71.58.100:admin:
...
```

Clearing Accounting Logs

To clear out the contents of the current log, use the **clear accounting log** command.

```
switch# clear accounting log
```

Configuring Cisco Access Control Servers

The Cisco Access Control Server (ACS) uses TACACS+ and RADIUS protocols to provide AAA services that ensure a secure environment.When using the AAA server, user management is normally done using Cisco ACS. Figure 33-3, Figure 33-4, Figure 33-5, and Figure 33-6 display ACS server user setup configurations for network-admin roles and multiple roles using either RADIUS or TACACS+.

Caution Cisco MDS SAN-OS does not support all numeric usernames, whether created with RADIUS or TACACS+, or created locally. Local users with all numeric names cannot be created. If an all numeric user name exists on an AAA server and is entered during login, the user is not logged in.

Figure 33-3 Configuring the network-admin Role When Using RADIUS

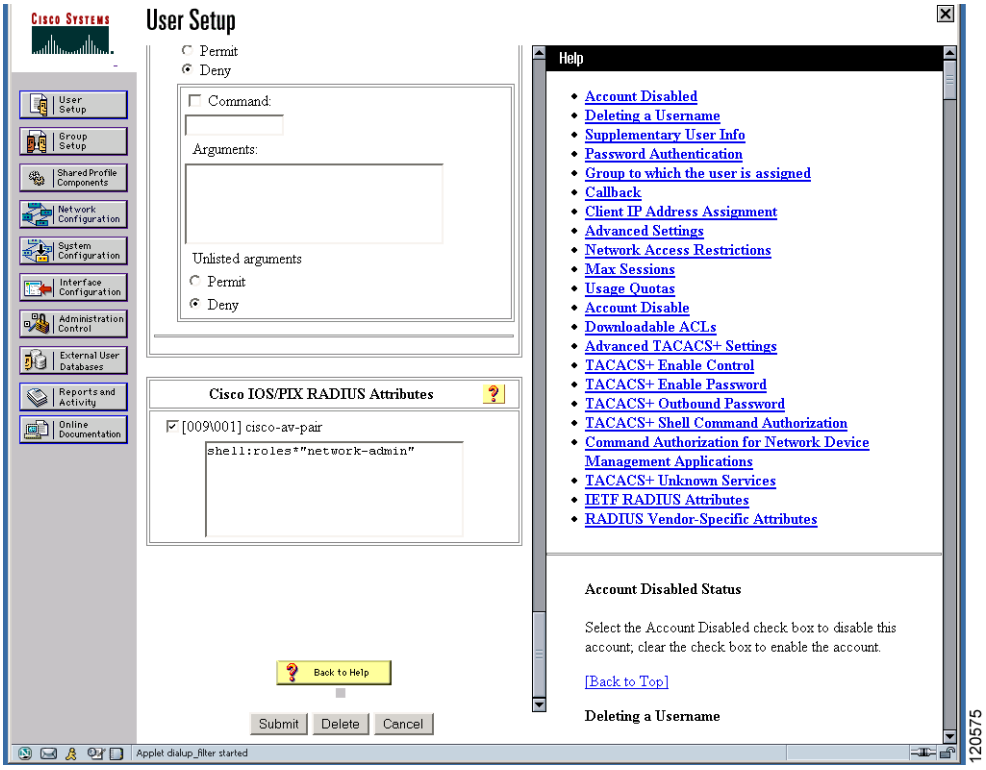

Figure 33-4 Configuring Multiple Roles with SNMPv3 Attributes When Using RADIUS

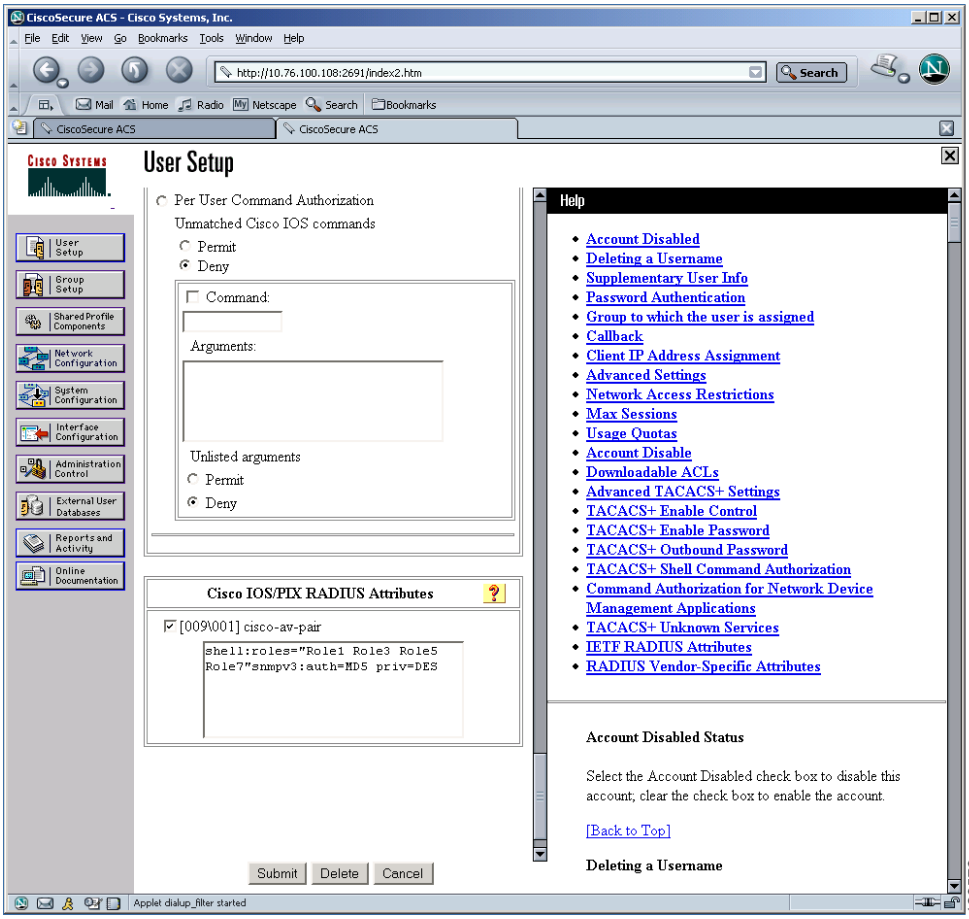

Figure 33-5 Configuring the network-admin Role with SNMPv3 Attributes When Using TACACS+

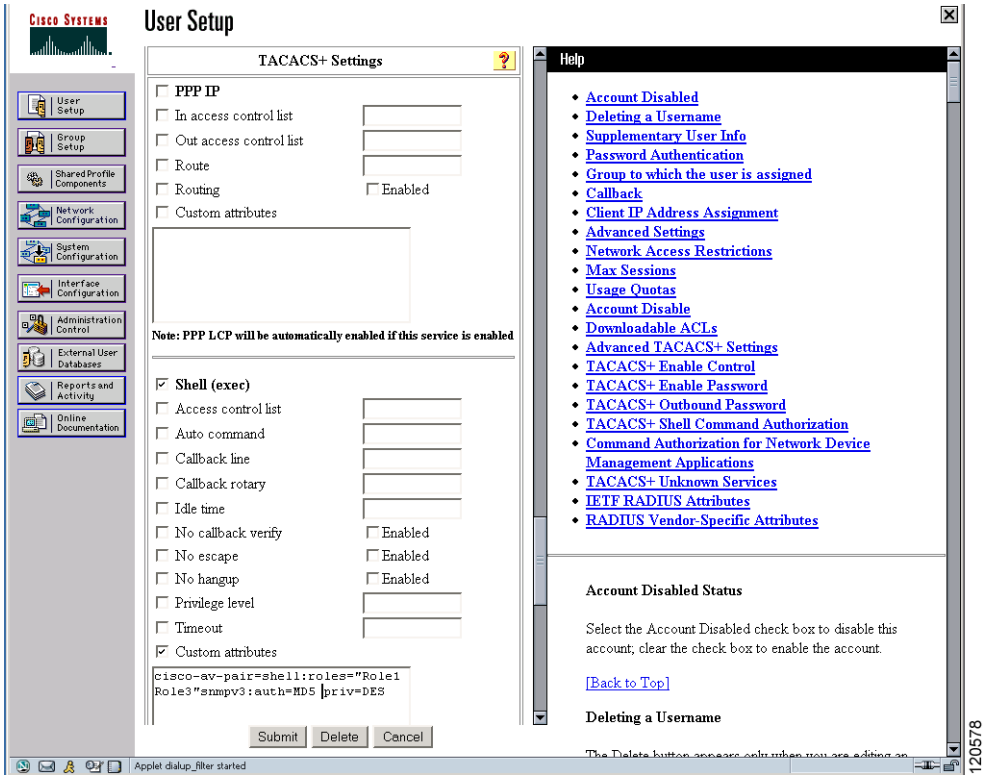

Figure 33-6 Configuring Multiple Roles with SNMPv3 Attributes When Using TACACS+

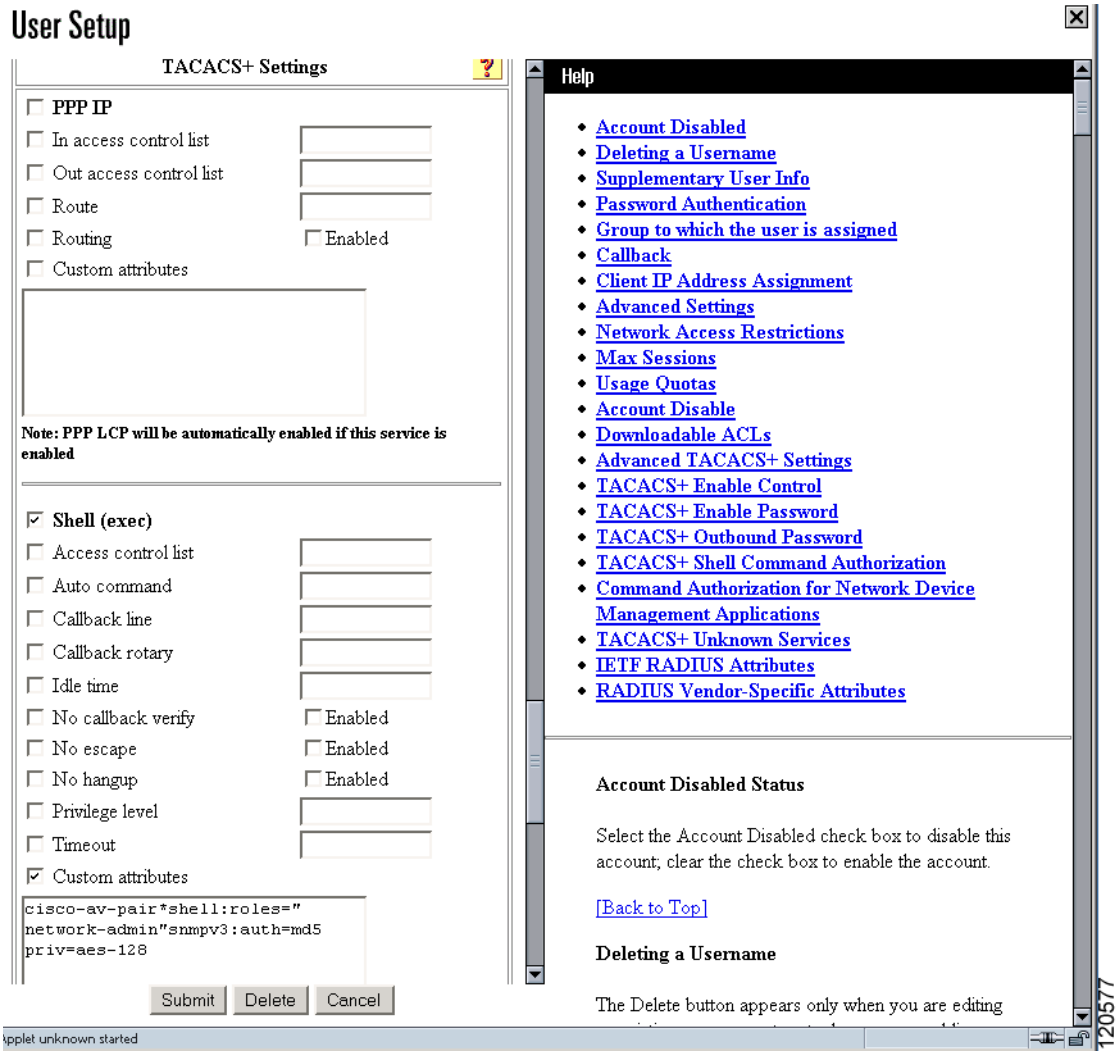

Default Settings

Table 33-2 lists the default settings for all switch security features in any switch.

Table 33-2 Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS switches	Network operator (network-operator)
AAA configuration services	Local
Authentication port	1821
Accounting port	1813
Preshared key communication	Clear text

Table 33-2 Default Switch Security Settings (continued)

Parameters	Default
RADIUS server timeout	1 (one) second
RADIUS server retries	Once
RADIUS server directed requests	Disabled
TACACS+	Disabled
TACACS+ servers	None configured
TACACS+ server timeout	5 seconds
TACACS+ server directed requests	Disabled
AAA server distribution	Disabled
Accounting log size	250 KB

CHAPTER 34

Configuring IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This chapter includes the following sections:

IPv4-ACL and IPv6-ACL Configuration Guidelines

Follow these guidelines when configuring IPv4-ACLs or IPv6-ACLs in any switch or director in the Cisco MDS 9000 Family:

- You can apply IPv4-ACLs or IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip If IPv4-ACLs or IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. See the "Gigabit Ethernet IPv4-ACL Guidelines" section on page 45-8 for guidelines on configuring IPv4-ACLs.

Caution Do not apply IPv4-ACLs or IPv6-ACLs to only one member of a PortChannel group. Apply IPv4-ACLs or IPv6-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv4-ACL or the IPv6-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TOS).

This section includes the following topics:

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

Note When configuring IPv4-ACLs or IPv6-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

The address information is required in each filter. It identifies the following details:

- Source: The address of the network or host from which the packet is being sent.
- Source-wildcard: The wildcard bits applied to the source.
- Destination: The number of the network or host to which the packet is being sent.
- Destination-wildcard: The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IPv4 address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IPv4 or IPv6 address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 requires an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.
- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. Table 34-1 displays the port numbers recognized by the Cisco SAN-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

Table 34-1 TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514

Table 34-1 TCP and UDP Port Numbers (continued)

Protocol	Port	Number
TCP[1]	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	wbem-http	5988
	wbem-https	5989

1. If the TCP connection is already established, use the **established** option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- The icmp-type: The ICMP message type is a number from 0 to 255.
- The icmp-code: The ICMP message code is a number from 0 to 255.

Table 34-2 displays the value for each ICMP type.

Table 34-2 ICMP Type Value

ICMP Type[1]	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

1. ICMP redirect packets are always rejected.

TOS Information

IP packets can be filtered based on the following optional TOS conditions:

- The TOS level: The level is specified by a number from 0 to 15.
- The TOS name: The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Configuring IPv4-ACLs or IPv6-ACLs

Traffic coming into the switch is compared to IPv4-ACL or IPv6-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv4-ACL or the IPv6-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv4-ACL or IPv6-ACL with only one deny entry has the effect of denying all traffic.

To configure an IPv4-ACL or an IPv6-ACL, you must complete the following tasks:

Step 1 Create an IPv4-ACL or an IPv6-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.

Note The filter entries are executed in sequential order. You can only add the entries to the end of the list. Take care to add the entries in the correct order.

Step 2 Apply the access filter to specified interfaces.

Creating IPv4-ACLs or IPv6-ACLs

To create an IPv4-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List1 permit ip any any`	Configures an IPv4-ACL called List1 and permits IP traffic from any source address to any destination address.
	`switch(config)# no ip access-list List1 permit ip any any`	Removes the IPv4-ACL called List1.
Step 3	`switch(config)# ip access-list List1 deny tcp any any`	Updates List1 to deny TCP traffic from any source address to any destination address.

To create an IPv6-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ipv6 access-list List1` `switch(config-ipv6-acl)#`	Configures an IPv6-ACL called List1 and enters IPv6-ACL configuration submode.
	`switch(config)# no ipv6 access-list List1`	Removes the IPv6-ACL called List1 and all its entries.

	Command	Purpose
Step 3	`switch(config-ipv6-acl)# permit ipv6 any any`	Adds an entry permitting IPv6 traffic from any source address to any destination address.
	`switch(config-ipv6-acl)# no permit ipv6 any any`	Removes an entry from the IPv6-ACL.
	`switch(config-ipv6-acl)# deny tcp any any`	Adds an entry to deny TCP traffic from any source address to any destination address.

To define an IPv4-ACL that restricts management access, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list restrict_mgmt permit ip 10.67.16.0 0.0.0.255 any`	Defines an entry in an IPv4-ACL named restrict_mgmt allowing all addresses in the 10.67.16.0/24 subnet.
Step 3	`switch(config)# ip access-list restrict_mgmt permit icmp any any eq 8`	Adds an entry to an IPv4-ACL named restrict_mgmt to allow any device to ping the MDS (icmp type 8).
Step 4	`switch(config)# ip access-list restrict_mgmt deny ip any any`	Explicitly blocks all other access to an access-list named restrict_mgmt.

To define an IPv6-ACL that restricts management access, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list RestrictMgmt` `switch(config-ipv6-acl)#`	Configures an IPv6-ACL called RestrictMgmt and enters IPv6-ACL configuration submode.
Step 3	`switch(config)# permit ipv6 2001:0DB8:800:200C::/64 any`	Defines an entry allowing all addresses in the 2001:0DB8:800:200C::/64 prefix.
Step 4	`switch(config)# permit icmp any any eq 8`	Adds an entry to allow any device to ping the MDS (ICMP type 8).
Step 5	`switch(config)# deny ipv6 any any`	Explicitly blocks all other IPv6 access.

To use the operand and port options for an IPv4-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any`	Denies TCP traffic from 1.2.3.0 through source port 5 to any destination.

To use the operand and port options for an IPv6-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List2 deny tcp 2001:0DB8:800:200C::/64 eq port 5 any`	Denies TCP traffic from 2001:0DB8:800:200C::/64 through source port 5 to any destination.

Adding IP Filters to an Existing IPv4-ACL or IPv6-ACL

After you create an IPv4-ACL or an IPv6-ACL, you can add subsequent IP filters at the end of the IPv4-ACL or the IPv6-ACL. You cannot insert filters in the middle of an IPv4-ACL or an IPv6-ACL. Each configured entry is automatically added to the end of a IPv4-ACL or a IPv6-ACL.

To add entries to an existing IPv4-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port telnet`	Permits TCP for Telnet traffic.
Step 3	`switch(config)# ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port http`	Permits TCP for HTTP traffic.
Step 4	`switch(config)# ip access-list List1 permit udp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0`	Permits UDP for all traffic.

To add entries to an existing IPv6-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ipv6 access-list List2` `switch(config-ipv6-acl)#`	Configures an IPv6-ACL and enters IPv6-ACL configuration submode.
Step 3	`switch(config-ipv6-acl)# permit ip 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64 eq 23`	Permits TCP for Telnet traffic.
Step 4	`switch(config-ipv6-acl)# permit tcp 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64 eq 143`	Permits TCP for HTTP traffic.
Step 5	`switch(config-ipv6-acl)# permit udp 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64`	Permits UDP for all traffic.

Removing IP Filters from an Existing IPv4-ACL or IPv6-ACL

To remove configured entries from an IPv4-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# no ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any`	Removes this entry from the IPv4-ACL (List2).
	`switch(config)# no ip access-list x3 deny ip any any`	Removes this entry from the IPv4-ACL (x3).
	`switch(config)# no ip access-list x3 permit ip any any`	Removes this entry from the IPv4-ACL (x3).

To remove configured entries from an IPv6-ACL, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ipv6 access-list List3` `switch(config-ipv6-acl)#`	Configures an IPv6-ACL and enters IPv6-ACL configuration submode.
Step 3	`switch(config-ipv6-acl)# no deny tcp 2001:0DB8:800:2010::/64 eq port 5 any`	Removes the TCP entry from the IPv6-ACL.
Step 4	`switch(config-ipv6-acl)# no deny ip any any`	Removes the IP entry from the IPv6-ACL.

Verifying the IPv4-ACL or IPv6-ACL Configuration

Use the **show ip access-list** command to view the contents of configured IPv4-ACLs. An IPv4-ACL can have one or more filters. (See Example 34-1.)

Example 34-1 Displays Filters Configured for an IPv4-ACL

```
switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)
```

Use the **show ipv6 access-list** command to view the contents of configured access filters. Each access filter can have several conditions. (See Example 34-2 and Example 34-3.)

Example 34-2 Displays Configured IPv6-ACLs

```
switch# show ipv6 access-list
Access List Name/Number          Filters IF   Status    Creation Time
-------------------------------- ------- ---- --------- -------------
abc                              3       7    active    Tue Jun 24 17:51:40 2003
x1                               3       1    active    Tue Jun 24 18:32:25 2003
x3                               0       1    not-ready Tue Jun 24 18:32:28 2003
```

Example 34-3 Displays a Summary of the Specified IPv6-ACL

```
switch# show ipv6 access-list abc
```

Reading the IP-ACL Log Dump

Use the **log-deny** option at the end of a filter condition to log information about packets that match dropped entries. The log output displays the ACL number, permit or deny status, and port information.

Note To capture these messages in a logging destination, you must configure severity level 7 for the kernel and ipacl facilities (see the "Facility Severity Levels" section on page 53-5) and severity level 7 for the logging destination: logfile (see the "Log Files" section on page 53-6), monitor (see the "Monitor Severity Level" section on page 53-5) or console (see the "Console Severity Level" section on page 53-4). For example:

```
switch# config t
switch(config)# logging level kernel 7
switch(config)# logging level ipacl 7
switch(config)# logging logfile message 7
```

For the input ACL, the log displays the raw MAC information. The keyword "MAC=" does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example is an input ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Applying an IP-ACL to an Interface

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to an interface on the switch. You can apply IP-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv4-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (see Figure 34-1).

Figure 34-1 Denying Traffic on the Inbound Interface

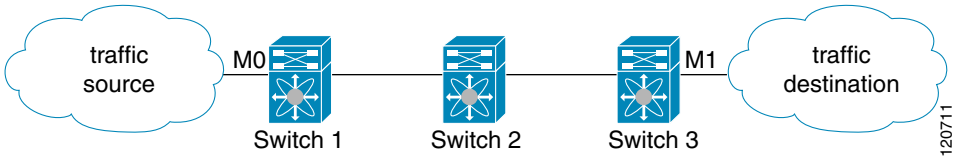

The **access-group** option controls access to an interface. Each interface can only be associated with one IP-ACL per direction. The ingress direction can have a different IP-ACL than the egress direction. The IP-ACL becomes active when applied to the interface.

Tip Create all conditions in an IP-ACL before applying it to the interface.

Caution If you apply an IP-ACL to an interface before creating it, all packets in that interface are dropped because the IP-ACL is empty.

The terms *in, out, source*, and *destination* are used as referenced by the switch:

- In—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).

Tip The IP-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.

Tip The IP-ACL applied to the interface for the egress traffic only affects local traffic.

To apply an IPv4-ACL to an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Configures a management interface (mgmt0).
Step 3	`switch(config-if)# ip access-group restrict_mgmt`	Applies an IPv4-ACL called restrict_mgmt for both the ingress and egress traffic (default).
	`switch(config-if)# no ip access-group NotRequired`	Removes the IPv4-ACL called NotRequired.

	Command	Purpose
Step 4	`switch(config-if)# ip access-group restrict_mgmt in`	Applies an IPv4-ACL called restrict_mgmt (if it does not already exist) for ingress traffic.
	`switch(config-if)# no ip access-group restrict_mgmt in`	Removes the IPv4-ACL called restrict_mgmt for ingress traffic.
	`switch(config-if)# ip access-group SampleName2 out`	Applies an IPv4-ACL called SampleName2 (if it does not already exist) for local egress traffic.
	`switch(config-if)# no ip access-group SampleName2 out`	Removes the IPv4-ACL called SampleName2 for local egress traffic.

To apply an IPv6-ACL to an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Configures a management interface (mgmt0).
Step 3	`switch(config-if)# ipv6 traffic-filter RestrictMgmt in`	Applies an IPv6-ACL called RestrictMgmt (if it does not already exist) for ingress traffic.
	`switch(config-if)# no ipv6 traffic-filter RestrictMgmt in`	Removes the IPv6-ACL called RestrictMgmt for ingress traffic.
	`switch(config-if)# ipv6 traffic-filter SampleName2 out`	Applies an IPv6-ACL called SampleName2 (if it does not already exist) for local egress traffic.
	`switch(config-if)# no ipv6 traffic-filter SampleName2 out`	Removes the IPv6-ACL called SampleName2 for local egress traffic.

Verifying Interface IP-ACL Configuration

Use the **show interface** command to display the IPv4-ACL configuration on an interface.

```
switch# show interface mgmt 0
mgmt0 is up
    Hardware is FastEthernet
    Address is 000c.30d9.fdbc
    Internet address is 172.22.31.113/24
    MTU 1500 bytes, BW 100 Mbps full Duplex
    ip access-group restrict_mgmt in
    35988 packets input, 3105539 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    2495 packets output, 430547 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Use the **show interface** command to display the IPv6-ACL configuration on an interface.

```
switch# show interface gigabitethernet 2/1
GigabitEthernet2/1 is up
    Hardware is GigabitEthernet, address is 000e.38c6.28b0
    Internet address is 10.1.1.10/24
    MTU 1500  bytes
    Port mode is IPS
    Speed is 1 Gbps
    Beacon is turned off
    Auto-Negotiation is turned on
    ip access-group RestrictMgmt
    5 minutes input rate 1208 bits/sec, 151 bytes/sec, 2 frames/sec
    5 minutes output rate 80 bits/sec, 10 bytes/sec, 0 frames/sec
    6232 packets input, 400990 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    503 packets output, 27054 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

IP-ACL Counter Cleanup

Use the **clear** command to clear the counters for a specified IPv4-ACL filter entry.

Note You cannot use this command to clear the counters for individual filters.

```
switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)

switch# clear ip access-list counters abc

switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (0 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (0 matches)
```

Use the **clear ipv6 access-list** command to clear the counters for all IPv6-ACLs.

```
switch# clear ipv6 access-list
```

Use the **clear ipv6 access-list** *name* command to clear the counters for a specified IPv6-ACL.

```
switch# clear ipv6 access-list List1
```

Note You cannot use this command to clear the counters for each individual filter.

.

CHAPTER 35

Configuring Certificate Authorities and Digital Certificates

Public Key Infrastructure (PKI) support provides the means for the Cisco MDS 9000 Family switches to obtain and use digital certificates for secure communication in the network. PKI support provides manageability and scalability for IPsec/IKE and SSH.

This chapter includes the following sections:

About CAs and Digital Certificates

This section provides information about certificate authorities (CAs) and digital certificates, and includes the following topics:

Purpose of CAs and Digital Certificates

CAs manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

Digital certificates link the digital signature to the sender. A digital certificate contains information to identify a user or device, such as the name, serial number, company, department, or IP address. It also contains a copy of the entity's public key. The certificate is itself signed by a CA, a third party that is explicitly trusted by the receiver to validate identities and to create digital certificates.

To validate the signature of the CA, the receiver must first know the CA's public key. Normally this process is handled out-of-band or through an operation done at installation. For instance, most web browsers are configured with the public keys of several CAs by default. The Internet Key Exchange (IKE), an essential component of IPsec, can use digital signatures to scalably authenticate peer devices before setting up security associations.

Trust Model, Trust Points, and Identity CAs

The trust model used in PKI support is hierarchical with multiple configurable trusted CAs. Each participating entity is configured with a list of CAs to be trusted so that the peer's certificate obtained during the security protocol exchanges can be verified, provided it has been issued by one of the locally trusted CAs. To accomplish this, CA's self signed root certificate (or certificate chain for a subordinate CA) is locally stored. The process of securely obtaining a trusted CA's root certificate (or the entire chain in the case of a subordinate CA) and storing it locally is called *CA authentication* and is a mandatory step in trusting a CA.

The information about a trusted CA that is locally configured is called the *trust point* and the CA itself is called a *trust point CA*. This information consists of CA certificate (or certificate chain in case of a subordinate CA) and the certificate revocation checking information.

The MDS switch can also enroll with a trust point to obtain an identity certificate (for example, for IPsec/IKE). This trust point is called an *identity CA*.

RSA Key-Pairs and Identity Certificates

You can generate one or more RSA key-pairs and associate each RSA key-pair with a trust point CA where the MDS switch intends to enroll to obtain an identity certificate. The MDS switch needs only one identity per CA, which consists of one key-pair and one identity certificate per CA.

Cisco MDS SAN-OS allows you to generate RSA key-pairs with a configurable key size (or modulus). The default key size is 512. You can also configure an RSA key-pair label. The default key label is the switch fully qualified domain name (FQDN).

The following list summarizes the relationship between trust points, RSA key-pairs, and identity certificates:

- A trust point corresponds to a specific CA that the MDS switch trusts for peer certificate verification for any application (such as IKE or SSH).
- An MDS switch can have many trust points and all applications on the switch can trust a peer certificate issued by any of the trust point CAs.
- A trust point is not restricted to a specific application.
- An MDS switch enrolls with the CA corresponding to the trust point to obtain an identity certificate. You can enroll your switch with multiple trust points thereby obtaining a separate identity certificate from each trust point. The identity certificates are used by applications depending upon the purposes specified in the certificate by the issuing CA. The purpose of a certificate is stored in the certificate as certificate extensions.
- When enrolling with a trust point, you must specify an RSA key-pair to be certified. This key-pair must be generated and associated to the trust point before generating the enrollment request. The association between the trust point, key-pair, and identity certificate is valid until it is explicitly removed by deleting the certificate, key-pair, or trust point.
- The subject name in the identity certificate is the fully qualified domain name for the MDS switch.
- You can generate one or more RSA key-pairs on a switch and each can be associated to one or more trust points. But no more than one key-pair can be associated to a trust point, which means only one identity certificate is allowed from a CA.
- If multiple identity certificates (each from a distinct CA) have been obtained, the certificate that an application selects to use in a security protocol exchange with a peer is application specific (see the "IPsec Digital Certificate Support" section on page 36-7 and the "SSH Authentication Using Digital Certificates" section on page 37-18).
- You do not need to designate one or more trust points for an application. Any application can use any certificate issued by any trust point as long as the certificate purpose satisfies the application requirements.
- You do not need more than one identity certificate from a trust point or more than one key-pair to be associated to a trust point. A CA certifies a given identity (name) only once and does not issue multiple certificates with the same subject name. If you need more than one identity certificate for a CA, then define another trust point for the same CA, associate another key-pair to it, and have it certified, provided CA allows multiple certificates with the same subject name.

Multiple Trusted CA Support

An MDS switch can be configured to trust multiple CAs by configuring multiple trust points and associating each with a distinct CA. With multiple trusted CAs, you do not have to enroll a switch with the specific CA that issued a certificate to a peer. Instead, you configure the switch with multiple trusted CAs that the peer trusts. A switch can then use a configured trusted CA to verify certificates offered by a peer that were not issued by the same CA defined in the identity of the switch.

Configuring multiple trusted CAs allows two or more switches enrolled under different domains (different CAs) to verify the identity of each other when using IKE to set up IPsec tunnels.

PKI Enrollment Support

Enrollment is the process of obtaining an identity certificate for the switch that is used for applications like IPsec/IKE or SSH. It occurs between the switch requesting the certificate and the certificate authority.

The PKI enrollment process for a switch involves the following steps:

1. Generate an RSA private and public key-pair on the switch.
2. Generate a certificate request in standard format and forward it to the CA.
3. Manual intervention at the CA server by the CA administrator may be required to approve the enrollment request, when it is received by the CA.
4. Receive the issued certificate back from the CA, signed with the CA's private key.
5. Write the certificate into a nonvolatile storage area on the switch (bootflash).

Manual Enrollment Using Cut-and-Paste Method

Cisco MDS SAN-OS supports certificate retrieval and enrollment using a manual cut-and-paste method. Cut-and-paste enrollment literally means you must cut and paste the certificate requests and resulting certificates between the switch and the CA, as follows:

1. Create an enrollment certificate request, which is displayed in base64-encoded text form.
2. Cut and paste the encoded certificate request text in an e-mail message or in a web form and send it to the CA.
3. Receive the issued certificate (in base64-encoded text form) from the CA in an e-mail message or in a web browser download.
4. Cut and paste the issued certificate to the switch using the certificate import facility.

Multiple RSA Key-Pair and Identity CA Support

Multiple identity CA support enables the switch to enroll with more than one trust point. This results in multiple identity certificates; each from a distinct CA. This allows the switch to participate in IPsec and other applications with many peers using certificates issued by appropriate CAs that are acceptable to those peers.

The multiple RSA key-pair support feature allows the switch to maintain a distinct key pair for each CA with which it is enrolled. Thus, it can match policy requirements for each CA without conflicting with the requirements specified by the other CAs, such as key length. The switch can generate multiple RSA key-pairs and associate each key-pair with a distinct trust point. Thereafter, when enrolling with a trust point, the associated key-pair is used to construct the certificate request.

Peer Certificate Verification

The PKI support on an MDS switch provides the means to verify peer certificates. The switch verifies certificates presented by peers during security exchanges pertaining to applications, such as IPsec/IKE and SSH. The applications verify the validity of the peer certificates presented to them. The peer certificate verification process involves the following steps:

- Verifies that the peer certificate is issued by one of the locally trusted CAs.
- Verifies that the peer certificate is valid (not expired) with respect to current time.
- Verifies that the peer certificate is not yet revoked by the issuing CA.

For revocation checking, two methods are supported: certificate revocation list (CRL) and Online Certificate Status Protocol (OCSP). A trust point uses one or both of these methods to verify that the peer certificate has not been revoked.

CRL Downloading, Caching, and Checking Support

Certificate revocation lists (CRLs) are maintained by CAs to give information of prematurely revoked certificates, and the CRLs are published in a repository. The download URL is made public and also specified in all issued certificates. A client verifying a peer's certificate should obtain the latest CRL from the issuing CA and use it to determine if the certificate has been revoked. A client can cache the CRLs of some or all of its trusted CAs locally and use them later if necessary until the CRLs expire.

Cisco MDS SAN-OS allows the manual configuration of pre-downloaded of CRLs for the trust points, and then caches them in the switch bootflash (cert-store). During the verification of a peer certificate by IPsec or SSH, the issuing CA's CRL is consulted only if the CRL has already been cached locally and the revocation checking is configured to use CRL. Otherwise, CRL checking is not performed and the certificate is considered to be not revoked if no other revocation checking methods are configured. This mode of CRL checking is called CRL optional.

OCSP Support

Online Certificate Status Protocol (OCSP) facilitates online certificate revocation checking. You can specify an OCSP URL for each trust point. Applications choose the revocation checking mechanisms in a specified order. The choices are CRL, OCSP, none, or a combination of these methods.

Import and Export Support for Certificates and Associated Key Pairs

As part of the CA authentication and enrollment process, the subordinate CA certificate (or certificate chain) and identity certificates can be imported in standard PEM (base64) format.

The complete identity information in a trust point can be exported to a file in the password-protected PKCS#12 standard format. It can be later imported to the same switch (for example, after a system crash) or to a replacement switch. The information in a PKCS#12 file consists of the RSA key-pair, the identity certificate, and the CA certificate (or chain).

Configuring CAs and Digital Certificates

This section describes the tasks you must perform to allow CAs and digital certificates your Cisco MDS switch device to interoperate. This section includes the following sections:

Configuring the Host Name and IP Domain Name

You must configure the host name and IP domain name of the switch if they are not already configured. This is required because switch FQDN is used as the subject in the identity certificate. Also, the switch FQDN is used as a default key label when none is specified during key-pair generation. For example, a certificate named SwitchA.example.com is based on a switch host name of SwitchA and a switch IP domain name of example.com.

Caution Changing the host name or IP domain name after generating the certificate can invalidate the certificate.

To configure the host name and IP domain name of the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# hostname SwitchA`	Configures the host name (SwitchA) of the switch.
Step 3	`SwitchA(config)# ip domain-name example.com`	Configures the IP domain name (example.com) of the switch.

Generating an RSA Key-Pair

RSA key-pairs are used to sign and/or encrypt and decrypt the security payload during security protocol exchanges for applications such as IKE/IPsec and SSH, and they are required before you can obtain a certificate for your switch.

To generate an RSA key-pair, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto key generate rsa`	Generates an RSA key-pair with the switch FQDN as the default label and 512 as the default modulus. By default, the key is not exportable. **Note** The security policy (or requirement) at the local site (MDS switch) and at the CA (where enrollment is planned) are considered in deciding the appropriate key modulus. **Note** The maximum number of key-pairs you can configure on a switch is 16.
	`switch(config)# crypto key generate rsa label SwitchA modulus 768`	Generates an RSA key-pair with the label SwitchA and modulus 768. Valid modulus values are 512, 768, 1024, 1536, and 2048. By default, the key is not exportable.
	`switch(config)# crypto key generate rsa exportable`	Generates an RSA key-pair with the switch FQDN as the default label and 512 as the default modulus. The key is exportable. **Caution** The exportability of a key-pair cannot be changed after key-pair generation. **Note** Only exportable key-pairs can be exported in PKCS#12 format.

Creating a Trust Point CA Association

To create a trust point CA association, follow these steps:

	Command	Purpose
Step 1	`switch(config)# crypto ca trustpoint admin-ca` `switch(config-trustpoint)#`	Declares a trust point CA that the switch should trust and enters trust point configuration submode. **Note** The maximum number of trust points you can declare on a switch is 16.
	`switch(config)# no crypto ca trustpoint admin-ca`	Removes the trust point CA.
Step 2	`switch(config-trustpoint)# enroll terminal`	Specifies manual cut-and-paste certificate enrollment (default). **Note** Manual cut-and-paste certificate enrollment is the only method supported for enrollment.
Step 3	`switch(config-trustpoint)# rsakeypair SwitchA`	Specifies the label of the RSA key-pair to be associated to this trust point for the purpose of enrollment. It was generated earlier in the "Generating an RSA Key-Pair" section on page 35-7. Only one RSA key-pair can be specified per CA.
	`switch(config-trustpoint)# no rsakeypair SwitchA`	Disassociates the RSA key-pair from the trust point (default).
Step 4	`switch(config-trustpoint)# end` `switch#`	Exits trust point configuration submode.
Step 5	`switch# copy running-config startup-config`	Copies the running configuration to the startup configuration to ensure the configuration is persistent across reboots.

Authenticating the CA

The configuration process of trusting a CA is complete only when the CA is authenticated to the MDS switch. The switch must authenticate the CA. It does this by obtaining the self-signed certificate of the CA in PEM format, which contains the public key of the CA. Because the certificate of the CA is self-signed (the CA signs its own certificate) the public key of the CA should be manually authenticated by contacting the CA administrator to compare the fingerprint of the CA certificate.

Note If the CA being authenticated is not a self-signed CA (that is, it is a subordinate CA to another CA, which itself may be a subordinate to yet another CA, and so on, finally ending in a self-signed CA), then the full list of the CA certificates of all the CAs in the certification chain needs to be input during the CA authentication step. This is called the *CA certificate chain* of the CA being authenticated. The maximum number of certificates in a CA certificate chain is 10.

To authenticate the certificate of the CA by cutting and pasting the certificate from an e-mail message or a website, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# crypto ca authenticate admin-ca`
	`input (cut & paste) CA certificate (chain) in PEM format;`
	`end the input with a line containing only END OF INPUT :`
	`-----BEGIN CERTIFICATE-----`
	`MIIC4jCCAoygAwIBAgIQBWDSiay0GZRPSRIljK0ZejANBgkqhkiG9w0BAQUFADCB`
	`kDEgMB4GCSqGSIb3DQEJARYRYW1hbmRrZUBjaXNjby5jb20xCzAJBgNVBAYTAklO`
	`MRIwEAYDVQQIEwlLYXJuYXRha2ExEjAQBgNVBAcTCUJhbmdhbG9yZTEOMAwGA1UE`
	`ChMFQ2lzY28xEzARBgNVBAsTCm5ldHN0b3JhZ2UxEjAQBgNVBAMTCUFwYXJuYSBD`
	`QTAeFw0wNTA1MDMyMjQ2MzdaFw0wNzA1MDMyMjU1MTdaMIGQMSAwHgYJKoZIhvcN`
	`AQkBFhFhbWFuZGtlQGNpc2NvLmNvbTELMAkGA1UEBhMCSU4xEjAQBgNVBAgTCUth`
	`cm5hdGFrYTESMBAGA1UEBxMJQmFuZ2Fsb3J1MQ4wDAYDVQQKEwVDaXNjbzETMBEG`
	`A1UECxMKbmV0c3RvcmFnZTESMBAGA1UEAxMJQXBhcm5hIENBMFwwDQYJKoZIhvcN`
	`AQEBBQADSwAwSAJBAMW/7b3+DXJPANBsIHHzluNccNM87ypyzwuoSNZXOMpeRXXI`
	`OzyBAgiXT2ASFuUOwQ1iDM8rO/41jf8RxvYKvysCAwEAAaOBvzCBvDALBgNVHQ8E`
	`BAMCAcYwDwYDVR0TAQH/BAUwAwEB/zAdBgNVHQ4EFgQUJyjyRoMbrCNMRU2OyRhQ`
	`GgsWbHEwawYDVR0fBGQwYjAuoCygKoYoaHR0cDovL3NzZS0wOC9DZXJ0RW5yb2xs`
	`L0FwYXJuYSUyMENBLmNybDAwoC6gLIYqZmlsZTovL1xcc3NlLTA4XENlcnRFbnJv`
	`bGxcQXBhcm5hJTIwQ0EuY3JsMBAGCSsGAQQBgjcVAQQDAgEAMA0GCSqGSIb3DQEB`
	`BQUAA0EAHv6UQ+8nE399Tww+KaGr0g0NIJaqNgLh0AFcT0rEyuyt/WYGPzksF9Ea`
	`NBG7E0oN66zex0EOEfG1Vs6mXp1//w==`
	`-----END CERTIFICATE-----`
	`END OF INPUT`
	`Fingerprint(s): MD5`
	`Fingerprint=65:84:9A:27:D5:71:03:33:9C:12:23:92:38:6F:78:12`
	
	`Do you accept this certificate? [yes/no]: y`
		Prompts you to cut and paste the certificate of the CA. Use the same name that you used when declaring the CA.
		Note The maximum number of trust points you can authenticate to a specific CA is 10.

Note For subordinate CA authentication, the full chain of CA certificates ending in a self-signed CA is required because the CA chain is needed for certificate verification as well as for PKCS#12 format export.

Configuring Certificate Revocation Checking Methods

During security exchanges with a client (for example, an IKE peer or SSH user), the MDS switch performs the certificate verification of the peer certificate sent by the client and the verification process may involve certificate revocation status checking.

You can use different methods for checking for revoked sender certificates. You can configure the switch to check the CRL downloaded from the CA (see the "Configuring a CRL" section on page 35-14), you can use OSCP if it is supported in your network, or both. Downloading the CRL and checking locally does not generate traffic in your network. However, certificates can be revoked between downloads and your switch would not be aware of the revocation. OCSP provides the means to check the current CRL on the CA. However, OCSP can generate network traffic that can impact network efficiency. Using both local CRL checking and OCSP provides the most secure method for checking for revoked certificates.

Note You must authenticate the CA before configuring certificate revocation checking.

To configure certificate revocation checking methods, follow these steps:

	Command	Purpose
Step 1	`switch(config)# crypto ca trustpoint admin-ca` `switch(config-trustpoint)#`	Declares a trust point CA that the switch should trust and enters trust point configuration submode.
Step 2	`switch(config-trustpoint)# ocsp url http://crlcheck.cisco.com`	Specifies the for OCSP to use to check for revoked certificates.
	`switch(config-trustpoint)# no ocsp url http://crlcheck.cisco.com`	Removes the URL for OCSP.
Step 3	`switch(config-trustpoint)# revocation-check oscp`	Specifies OCSP as the revocation checking method to be employed during verification of peer certificates issued by the same CA as that of this trust point. **Note** The OSCP URL must be configured before specifying OSCP as a revocation checking method.
	`switch(config-trustpoint)# revocation-check crl`	Specifies CRL (default) as the revocation checking method to be employed during verification of peer certificates issued by the same CA as that of this trust point.
	`switch(config-trustpoint)# revocation-check crl oscp`	Specifies CRL as the first revocation checking method and OCSP as the next method. If the CRL method fails (for example, due to the CRL is not found or has expired) to be used during verification of peer certificates issued by the same CA as that of this trust point, then OSCP is used. **Note** The OSCP URL must be configured before specifying OSCP as a revocation checking method.
	`switch(config-trustpoint)# revocation-check none`	Does not check for revoked certificates.
	`switch(config-trustpoint)# no revocation-check`	Reverts to default method.

Generating Certificate Requests

You must generate a request to obtain identity certificates from the associated trust point CA for each of your switch's RSA key-pairs. You must then cut and paste the displayed request into an e-mail message or in a website form for the CA.

To generate a request for signed certificates from the CA, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# crypto ca enroll admin-ca`
	`Create the certificate request ..`
	`Create a challenge password. You will need to verbally provide this password to the CA Administrator in order to revoke your certificate.`
	`For security reasons your password will not be saved in the configuration.`
	`Please make a note of it.`
	`Password:nbv123`
	`The subject name in the certificate will be: Vegas-1.cisco.com`
	`Include the switch serial number in the subject name? [yes/no]: no`
	`Include an IP address in the subject name [yes/no]: yes`
	`ip address:172.22.31.162`
	`The certificate request will be displayed...`
	`-----BEGIN CERTIFICATE REQUEST-----`
	`MIIBqzCCARQCAQAwHDEaMBgGA1UEAxMRVmVnYXMtMS5jaXNjby5jb20wgZ8wDQYJ`
	`KoZIhvcNAQEBBQADgY0AMIGJAoGBAL8Y1UAJ2NC7jUJ1DVaSMqNIgJ2kt8rl4lKY`
	`0JC6ManNy4qxk8VeMXZSiLJ4JgTzKWdxbLDkTTysnjuCXGvjb+wj0hEhv/y51T9y`
	`P2NJJ8ornqShrvFZgC7ysN/PyMwKcgzhbVpj+rargZvHtGJ91XTq4WoVkSCzXv8S`
	`VqyH0vEvAgMBAAGgTzAVBgkqhkiG9w0BCQcxCBMGbmJ2MTIzMDYGCSqGSIb3DQEJ`
	`DjEpMCcwJQYDVR0RAQH/BBswGYIRVmVnYXMtMS5jaXNjby5jb22HBKwWH6IwDQYJ`
	`KoZIhvcNAQEEBQADgYEAkT60KER6Qo8nj0sDXZVHSfJZh6K6JtDz3Gkd99GlFWgt`
	`PftrNcWUE/pw6HayfQl2T3ecgNwel2d15133YBF2bktExiI6Ul88nTOjglXMjja8`
	`8a23bNDpNsM8rklwA6hWkrVL8NUZEFJxqbjfngPNTZacJCUS6ZqKCMetbKytUx0=`
	`-----END CERTIFICATE REQUEST-----`
		Generates a certificate request for an authenticated CA.
		Note The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.

Installing Identity Certificates

You receive the identity certificate from the CA by e-mail or through a web browser in base64 encoded text form. You must install the identity certificate from the CA by cutting and pasting the encoded text using the CLI import facility.

To install an identity certificate received from the CA by e-mail or through a web browser, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# crypto ca import admin-ca certificate`
	`input (cut & paste) certificate in PEM format:`
	`-----BEGIN CERTIFICATE-----`
	`MIIEADCCA6qgAwIBAgIKCjOOoQAAAAAAdDANBgkqhkiG9w0BAQUFADCBkDEgMB4G`
	`CSqGSIb3DQEJARYRYW1hbmRrZUBjaXNjby5jb20xCzAJBgNVBAYTAklOMRIwEAYD`
	`VQQIEwlLYXJuYXRha2ExEjAQBgNVBAcTCUJhbmdhbG9yZTEOMAwGA1UEChMFQ2lz`
	`Y28xEzARBgNVBAsTCm5ldHN0b3JhZ2UxEjAQBgNVBAMTCUFwYXJuYSBDQTAeFw0w`
	`NTExMTIwMzAyNDBaFw0wNjExMTIwMzEyNDBaMBwxGjAYBgNVBAMTEVZlZ2FzLTEu`
	`Y2lzY28uY29tMIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQC/GNVACdjQu41C`
	`dQ1WkjKjSICdpLfK5eJSmNCQujGpzcuKsZPFXjF2UoiyeCYE8ylncWyw5E08rJ47`
	`glxr42/sI9IRIb/8udU/cj9jSSfKK56koa7xWYAu8rDfz8jMCnIM4W1aY/q2q4Gb`
	`x7RifdV06uFqFZEgs17/Elash9LxLwIDAQABo4ICEzCCAg8wJQYDVR0RAQH/BBsw`
	`GYIRVmVnYXMtMS5jaXNjby5jb22HBKwWH6IwHQYDVR0OBBYEFKCLi+2sspWEfgrR`
	`bhWmlVyo9jngMIHMBgNVHSMEgcQwgcGAFCco8kaDG6wjTEVNjskYUBoLFmxxoYGW`
	`pIGTMIGQMSAwHgYJKoZIhvcNAQkBFhFhbWFuZGtlQGNpc2NvLmNvbTELMAkGA1UE`
	`BhMCSU4xEjAQBgNVBAgTCUthcm5hdGFrYTESMBAGA1UEBxMJQmFuZ2Fsb3JlMQ4w`
	`DAYDVQQKEwVDaXNjbzETMBEGA1UECxMKbmV0c3RvcmFnZTESMBAGA1UEAxMJQXBh`
	`cm5hIENBghAFYNKJrLQZlE9JEiWMrRl6MGsGA1UdHwRkMGIwLqAsoCqGKGh0dHA6`
	`Ly9zc2UtMDgvQ2VydEVucm9sbC9BcGFybmElMjBDQS5jcmwwMKAuoCyGKmZpbGU6`
	`Ly9cXHNzZS0wOFxDZXJ0RW5yb2xsXEFwYXJuYSUyMENBLmNybDCBigYIKwYBBQUH`
	`AQEEfjB8MDsGCCsGAQUFBzAChi9odHRwOi8vc3NlLTA4L0NlcnRFbnJvbGwvc3Nl`
	`LTA4X0FwYXJuYSUyMENBLmNydDA9BggrBgEFBQcwAoYxZmlsZTovL1xcc3NlLTA4`
	`XENlcnRFbnJvbGxcc3NlLTA4X0FwYXJuYSUyMENBLmNydDANBgkqhkiG9w0BAQUF`
	`AANBADbGBGsbe7GNLh9xeOTWBNbm24U69ZSuDDcOcUZUUTgrpnTqVpPyejtsyf1w`
	`E36cIZu4WsExREqxbTk8ycx7V5o=`
	`-----END CERTIFICATE-----`
		Prompts you to cut and paste the identity certificate for the CA named admin-ca.
		Note The maximum number of identify certificates you can configure on a switch is 16.

Ensuring Trust Point Configurations Persist Across Reboots

The trust point configuration is a normal Cisco SAN-OS configuration that persists across system reboots only if you copy it explicitly to the startup configuration. The certificates, key-pairs, and CRL associated with a trust point are automatically persistent if you have already copied the trust point configuration in the startup configuration. Conversely, if the trust point configuration is not copied to the startup configuration, the certificates, key-pairs, and CRL associated with it are not persistent since they require the corresponding trust point configuration after a reboot. Always copy the running configuration to the startup configuration to ensure the that the configured certificates, key-pairs, and CRLs are persistent. Also, save the running configuration after deleting a certificate or key-pair to ensure the deletions permanent.

The certificates and CRL associated with a trust point automatically become persistent when imported (that is, without an explicitly copying to the startup configuration) if the specific trust point is already saved in startup configuration.

We also recommend that you create a password protected backup of the identity certificates nd save it to an external server (see the "Exporting and Importing Identity Information in PKCS#12 Format" section on page 35-13).

Note Copying the configuration to an external server does include the certificates and key-pairs.

Monitoring and Maintaining CA and Certificates Configuration

The tasks in the section are optional. This section includes the following topics:

Exporting and Importing Identity Information in PKCS#12 Format

You can export the identity certificate along with the RSA key-pair and CA certificate (or the entire chain in the case of a subordinate CA) of a trust point to a PKCS#12 file for backup purposes. You can later import the certificate and RSA key-pair to recover from a system crash on your switch or when you replace the supervisor modules.

Note Only bootflash:*filename* syntax is supported when specifying the export and import URL.

To export a certificate and key-pair to a PKCS#12-formatted file, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ca export admin-ca pkcs12 bootflash:adminid.p12 nbv123`	Exports the identity certificate and associated key-pair and CA certificates for trust point admin-ca to the file bootflash:adminid.p12 in PKCS#12 format, protected using password nbv123.
Step 3	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 4	`switch# copy bootflash:adminid.p12 tftp:adminid.p12`	Copies the PKCS#12 format file to a TFTP server.

To import a certificate and key-pair from a PKCS#12-formatted file, follow these steps:

	Command	Purpose
Step 1	`switch# copy tftp:adminid.p12 bootflash:adminid.p12`	Copies the PKCS#12 format file from a TFTP server.
Step 2	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 3	`switch(config)# crypto ca import admin-ca pkcs12 bootflash:adminid.p12 nbv123`	Imports the identity certificate and associated key-pair and CA certificates for trust point admin-ca from the file bootflash:adminid.p12 in PKCS#12 format, protected using password nbv123.

Note :The trust point must be empty (with no RSA key-pair associated with it and no CA is associated with it using CA authentication) for the PKCS#12 file import to succeed.

Configuring a CRL

To import the CRL from a file to a trust point, follow these steps:

	Command	Purpose
Step 1	`switch# copy tftp:adminca.crl bootflash:adminca.crl`	Download the CRL.
Step 2	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 3	`switch(config)# crypto ca crl request admin-ca bootflash:adminca.crl`	Configures or replaces the current CRL with the one specified in the file.

Deleting Certificates from the CA Configuration

You can delete the identity certificates and CA certificates that are configured in a trust point. You must first delete the identity certificate, followed by the CA certificates. Then after deleting the identity certificate, you can disassociate the RSA key-pair from a trust point. The certificate deletion is necessary to remove expired or revoked certificates, certificates whose key-pairs are compromised (or suspected to be compromised) or CAs that are no longer trusted.

To delete the CA certificate (or the entire chain in the case of a subordinate CA) from a trust point, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ca trustpoint myCA`	Enters trustpoint configuration submode.
Step 3	`switch(config-trustpoint)# delete ca-certificate`	Deletes the CA certificate or certificate chain.
Step 4	`switch(config-trustpoint)# delete certificate`	Deletes the identity certificate.
	`switch(config-trustpoint)# delete certificate force`	Forces the deletion of the identity certificate. **Note** If the identity certificate being deleted is the last-most or only identity certificate in the device, you must use the **force** option to delete it. This ensures that the administrator does not mistakenly delete the last-most or only identity certificate and leave the applications (such as IKE and SSH) without a certificate to use.
Step 5	`switch(config-trustpoint)# end` `switch#`	Returns to EXEC mode.
Step 6	`switch# copy running-config startup-config`	Copies the running configuration to the startup configuration to ensure the configuration is persistent across reboots.

Deleting RSA Key-Pairs from Your Switch

Under certain circumstances you may want to delete your switch's RSA key-pairs. For example, if you believe the RSA key-pairs were compromised in some way and should no longer be used, you should delete the key-pairs.

To delete RSA key-pairs from your switch, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto key zeroize rsa MyKey`	Deletes the RSA key-pair whose label is MyKey.
Step 3	`switch(config)# end` `switch#`	Returns to EXEC mode.
Step 4	`switch# copy running-config startup-config`	Copies the running configuration to the startup configuration to ensure the configuration is persistent across reboots.

Note After you delete RSA key-pairs from a switch, ask the CA administrator to revoke your switch's certificates at the CA. You must supply the challenge password you created when you originally requested the certificates. See "Generating Certificate Requests" section on page 35-10.

Displaying Key-Pair and CA Information

To view key-pair and CA information, use the following commands in EXEC mode:

Command	Purpose
`switch# show crypto key mypubkey rsa`	Displays information about the switch's RSA public keys.
`switch# show crypto ca certificates`	Displays information on CA and identity certificates.
`switch# show crypto ca crl`	Displays information about CA CRLs.
`switch# show crypto ca trustpoints`	Displays information about CA trust points.

Example Configurations

This section shows an example of the tasks you can use to configure certificates and CRLs on the Cisco MDS 9000 Family switches using the Microsoft Windows Certificate server.

This section includes the following topics:

Configuring Certificates on the MDS Switch

To configure certificates on an MDS switch, follow these steps:

Step 1 Configure the switch FQDN.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# switchname Vegas-1
Vegas-1(config)#
```

Step 2 Configure the DNS domain name for the switch.

```
Vegas-1(config)# ip domain-name cisco.com
Vegas-1(config)#
```

Step 3 Create a trust point.

```
Vegas-1(config)# crypto ca trustpoint myCA
Vegas-1(config-trustpoint)# exit
Vegas-1(config)# do show crypto ca trustpoints
trustpoint: myCA; key:
revokation methods:  crl
Vegas-1(config)#
```

Step 4 Create an RSA key-pair for the switch.

```
Vegas-1(config)# crypto key generate rsa label myKey exportable modulus 1024
Vegas-1(config)# do show crypto key mypubkey rsa
key label: myKey
key size: 1024
exportable: yes

Vegas-1(config)#
```

Step 5 Associate the RSA key-pair to the trust point.

```
Vegas-1(config)# crypto ca trustpoint myCA
Vegas-1(config-trustpoint)# rsakeypair myKey
Vegas-1(config-trustpoint)# exit
Vegas-1(config)# do show crypto ca trustpoints
trustpoint: myCA; key: myKey
revokation methods:  crl
Vegas-1(config)#
```

Step 6 Download the CA certificate from the Microsoft Certificate Service web interface (see the "Downloading a CA Certificate" section on page 35-19)

Step 7 Authenticate the CA that you want to enroll to the trust point.

```
Vegas-1(config)# crypto ca authenticate myCA
input (cut & paste) CA certificate (chain) in PEM format;
end the input with a line containing only END OF INPUT :
-----BEGIN CERTIFICATE-----
MIIC4jCCAoygAwIBAgIQBWDSiay0GZRPSRIljK0ZejANBgkqhkiG9w0BAQUFADCB
kDEgMB4GCSqGSIb3DQEJARYRYW1hbmRrZUBjaXNjby5jb20xCzAJBgNVBAYTAklO
MRIwEAYDVQQIEwlLYXJuYXRha2ExEjAQBgNVBAcTCUJhbmdhbG9yZTEOMAwGA1UE
ChMFQ2lzY28xEzARBgNVBAsTCm5ldHN0b3JhZ2UxEjAQBgNVBAMTCUFwYXJuYSBD
QTAeFw0wNTA1MDMyMjQ2MzdaFw0wNzA1MDMyMjU1MTdaMIGQMSAwHgYJKoZIhvcN
AQkBFhFhbWFuZGtlQGNpc2NvLmNvbTELMAkGA1UEBhMCSU4xEjAQBgNVBAgTCUth
cm5hdGFrYTESMBAGA1UEBxMJQmFuZ2Fsb3JlMQ4wDAYDVQQKEwVDaXNjbzETMBEG
A1UECxMKbmV0c3RvcmFnZTESMBAGA1UEAxMJQXBhcm5hIENBMFwwDQYJKoZIhvcN
AQEBBQADSwAwSAJBAMW/7b3+DXJPANBsIHHzluNccNM87ypyzwuoSNZXOMpeRXXI
OzyBAgiXT2ASFuUOwQ1iDM8rO/41jf8RxvYKvysCAwEAAaOBvzCBvDALBgNVHQ8E
BAMCAcYwDwYDVR0TAQH/BAUwAwEB/zAdBgNVHQ4EFgQUJyjyRoMbrCNMRU2OyRhQ
GgsWbHEwawYDVR0fBGQwYjAuoCygKoYoaHR0cDovL3NzZS0wOC9DZXJ0RW5yb2xs
L0FwYXJuYSUyMENBLmNybDAwoC6gLIYqZmlsZTovL1xcc3NlLTA4XENlcnRFbnJv
bGxcQXBhcm5hJTIwQ0EuY3JsMBAGCSsGAQQBgjcVAQQDAgEAMA0GCSqGSIb3DQEB
BQUAA0EAHv6UQ+8nE399Tww+KaGr0g0NIJaqNgLh0AFcT0rEyuyt/WYGPzksF9Ea
NBG7E0oN66zex0EOEfG1Vs6mXp1//w==
-----END CERTIFICATE-----
 END OF INPUT
Fingerprint(s): MD5 Fingerprint=65:84:9A:27:D5:71:03:33:9C:12:23:92:38:6F:78:12



Do you accept this certificate? [yes/no]:y
Vegas-1(config)#

Vegas-1(config)# do show crypto ca certificates
Trustpoint: myCA
CA certificate 0:
subject= /emailAddress=admin@yourcompany.com/C=IN/ST=Karnataka/L=Bangalore/O=Yourcompany/O
U=netstorage/CN=Aparna CA
issuer= /emailAddress=admin@yourcompany.com/C=IN/ST=Karnataka/L=Bangalore/O=Yourcompany/OU
=netstorage/CN=Aparna CA
serial=0560D289ACB419944F4912258CAD197A
notBefore=May  3 22:46:37 2005 GMT
notAfter=May  3 22:55:17 2007 GMT
MD5 Fingerprint=65:84:9A:27:D5:71:03:33:9C:12:23:92:38:6F:78:12
purposes: sslserver sslclient ike
```

Step 8 Generate a request certificate to use to enroll with a trust point.

```
Vegas-1(config)# crypto ca enroll myCA
 Create the certificate request ..
 Create a challenge password. You will need to verbally provide this
  password to the CA Administrator in order to revoke your certificate.
  For security reasons your password will not be saved in the configuration.
  Please make a note of it.
  Password:nbv123
 The subject name in the certificate will be: Vegas-1.cisco.com
 Include the switch serial number in the subject name? [yes/no]:no
 Include an IP address in the subject name [yes/no]:yes
```

```
ip address:10.10.1.1
 The certificate request will be displayed...
-----BEGIN CERTIFICATE REQUEST-----
MIIBqzCCARQCAQAwHDEaMBgGA1UEAxMRVmVnYXMtMS5jaXNjby5jb20wgZ8wDQYJ
KoZIhvcNAQEBBQADgY0AMIGJAoGBAL8Y1UAJ2NC7jUJ1DVaSMqNIgJ2kt8rl4lKY
0JC6ManNy4qxk8VeMXZSiLJ4JgTzKWdxbLDkTTysnjuCXGvjb+wj0hEhv/y51T9y
P2NJJ8ornqShrvFZgC7ysN/PyMwKcgzhbVpj+rargZvHtGJ91XTq4WoVkSCzXv8S
VqyH0vEvAgMBAAGgTzAVBgkqhkiG9w0BCQcxCBMGbmJ2MTIzMDYGCSqGSIb3DQEJ
DjEpMCcwJQYDVR0RAQH/BBswGYIRVmVnYXMtMS5jaXNjby5jb22HBKwWH6IwDQYJ
KoZIhvcNAQEEBQADgYEAkT60KER6Qo8nj0sDXZVHSfJZh6K6JtDz3Gkd99GlFWgt
PftrNcWUE/pw6HayfQl2T3ecgNwel2d15133YBF2bktExiI6Ul88nTOjglXMjja8
8a23bNDpNsM8rklwA6hWkrVL8NUZEFJxqbjfngPNTZacJCUS6ZqKCMetbKytUx0=
-----END CERTIFICATE REQUEST-----

Vegas-1(config)#
```

Step 9 Request an identity certificate from the Microsoft Certificate Service web interface (see the "Requesting an Identity Certificate" section on page 35-23).

Step 10 Import the identity certificate.

```
Vegas-1(config)# crypto ca import myCA certificate
input (cut & paste) certificate in PEM format:
-----BEGIN CERTIFICATE-----
MIIEADCCA6qgAwIBAgIKCjOOoQAAAAAAdDANBgkqhkiG9w0BAQUFADCBkDEgMB4G
CSqGSIb3DQEJARYRYW1hbmRrZUBjaXNjby5jb20xCzAJBgNVBAYTAklOMRIwEAYD
VQQIEwlLYXJuYXRha2ExEjAQBgNVBAcTCUJhbmdhbG9yZTEOMAwGA1UEChMFQ2lz
Y28xEzARBgNVBAsTCm5ldHN0b3JhZ2UxEjAQBgNVBAMTCUFwYXJuYSBDQTAeFw0w
NTExMTIwMzAyNDBaFw0wNjExMTIwMzEyNDBaMBwxGjAYBgNVBAMTEVZlZ2FzLTEu
Y2lzY28uY29tMIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQC/GNVACdjQu41C
dQ1WkjKjSICdpLfK5eJSmNCQujGpzcuKsZPFXjF2UoiyeCYE8ylncWyw5E08rJ47
glxr42/sI9IRIb/8udU/cj9jSSfKK56koa7xWYAu8rDfz8jMCnIM4W1aY/q2q4Gb
x7RifdV06uFqFZEgs17/Elash9LxLwIDAQABo4ICEzCCAg8wJQYDVR0RAQH/BBsw
GYIRVmVnYXMtMS5jaXNjby5jb22HBKwWH6IwHQYDVR0OBBYEFKCLi+2sspWEfgrR
bhWmlVyo9jngMIHMBgNVHSMEgcQwgcGAFCco8kaDG6wjTEVNjskYUBoLFmxxoYGW
pIGTMIGQMSAwHgYJKoZIhvcNAQkBFhFhbWFuZGtlQGNpc2NvLmNvbTELMAkGA1UE
BhMCSU4xEjAQBgNVBAgTCUthcm5hdGFrYTESMBAGA1UEBxMJQmFuZ2Fsb3JlMQ4w
DAYDVQQKEwVDaXNjbzETMBEGA1UECxMKbmV0c3RvcmFnZTESMBAGA1UEAxMJQXBh
cm5hIENBghAFYNKJrLQZ1E9JEiWMrRl6MGsGA1UdHwRkMGIwLqAsoCqGKGh0dHA6
Ly9zc2UtMDgvQ2VydEVucm9sbC9BcGFybmElMjBDQS5jcmwwMKAuoCyGKmZpbGU6
Ly9cXHNzZS0wOFxDZXJ0RW5yb2xsXEFwYXJuYSUyMENBLmNybDCBigYIKwYBBQUH
AQEEfjB8MDsGCCsGAQUFBzAChi9odHRwOi8vc3NlLTA4L0NlcnRFbnJvbGwvc3Nl
LTA4X0FwYXJuYSUyMENBLmNydDA9BggrBgEFBQcwAoYxZmlsZTovL1xcc3NlLTA4
XENlcnRFbnJvbGxcc3NlLTA4X0FwYXJuYSUyMENBLmNydDANBgkqhkiG9w0BAQUF
AANBADbGBGsbe7GNLh9xeOTWBNbm24U69ZSuDDcOcUZUUTgrpnTqVpPyejtsyf1w
E36cIZu4WsExREqxbTk8ycx7V5o=
-----END CERTIFICATE-----
Vegas-1(config)# exit
Vegas-1#
```

Step 11 Verify the certificate configuration.

```
Vegas-1# show crypto ca certificates
Trustpoint: myCA
certificate:
subject= /CN=Vegas-1.cisco.com
issuer= /emailAddress=admin@yourcompany.com/C=IN/ST=Karnataka/L=Bangalore/O=Cisco/OU
=netstorage/CN=Aparna CA
serial=0A338EA1000000000074
notBefore=Nov 12 03:02:40 2005 GMT
notAfter=Nov 12 03:12:40 2006 GMT
MD5 Fingerprint=3D:33:62:3D:B4:D0:87:A0:70:DE:A3:87:B3:4E:24:BF
purposes: sslserver sslclient ike
```

```
CA certificate 0:
subject= /emailAddress=admin@yourcompany.com/C=IN/ST=Karnataka/L=Bangalore/O=Yourcompany/O
U=netstorage/CN=Aparna CA
issuer= /emailAddress=admin@yourcompany.com/C=IN/ST=Karnataka/L=Bangalore/O=Yourcompany/OU
=netstorage/CN=Aparna CA
serial=0560D289ACB419944F4912258CAD197A
notBefore=May  3 22:46:37 2005 GMT
notAfter=May  3 22:55:17 2007 GMT
MD5 Fingerprint=65:84:9A:27:D5:71:03:33:9C:12:23:92:38:6F:78:12
purposes: sslserver sslclient ike
```

Step 12 Save the certificate configuration to the startup configuration.

```
Vegas-1# copy running-config startup-config
```

Downloading a CA Certificate

To download a CA certificate from the Microsoft Certificate Services web interface, follow these steps:

Step 1 Select the **Retrieve the CA certificate or certificate revocation task** radio button in the Microsoft Certificate Services web interface and click the **Next** button.

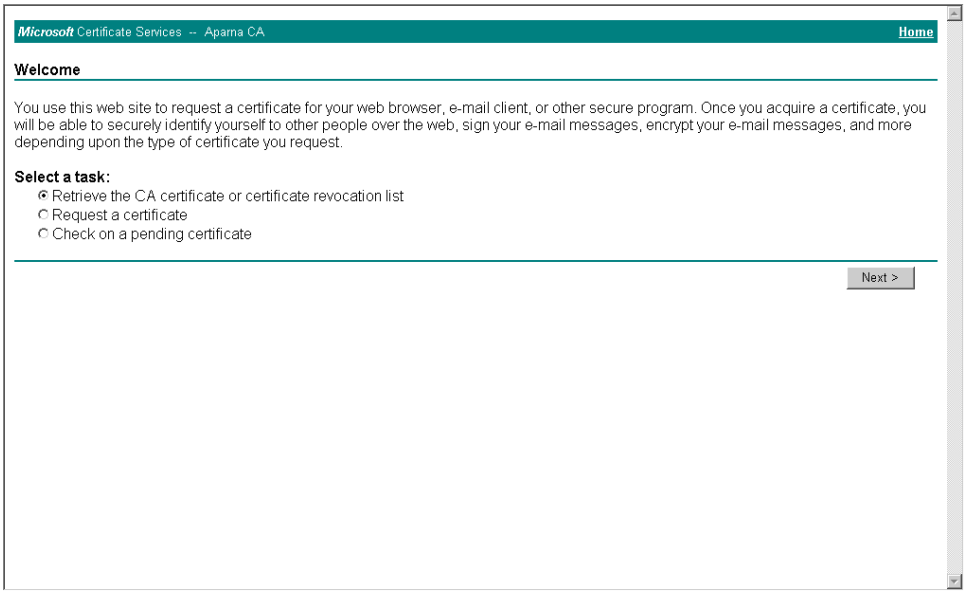

Step 2 Select the CA certificate file to download from the displayed list. Click the **Base 64 encoded** radio button, and click the **Download CA certificate** link.

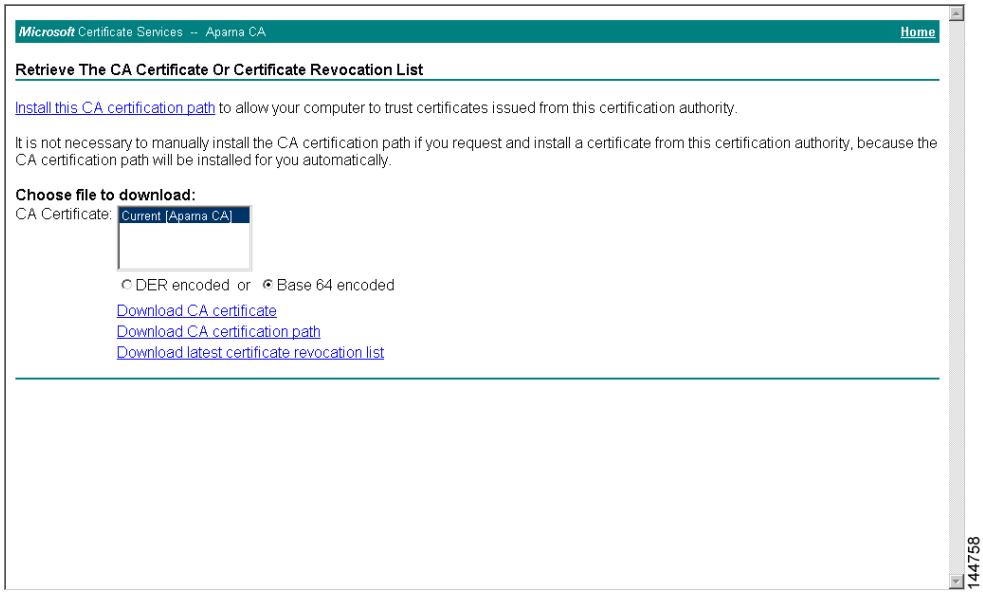

Step 3 Click the **Open** button in the File Download dialog box.

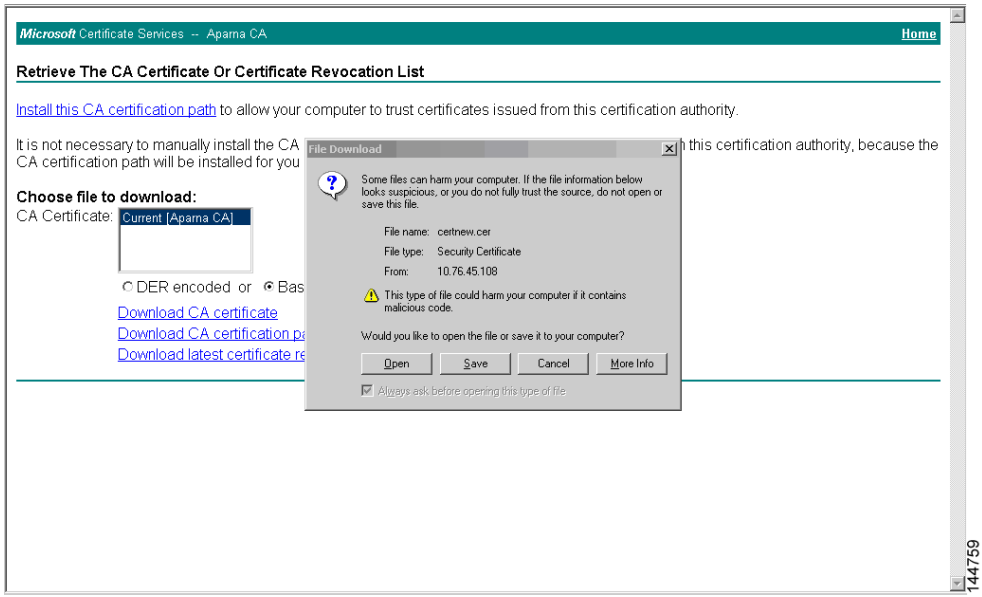

Step 4 Click the **Copy to File** button in the Certificate dialog box and click **OK**.

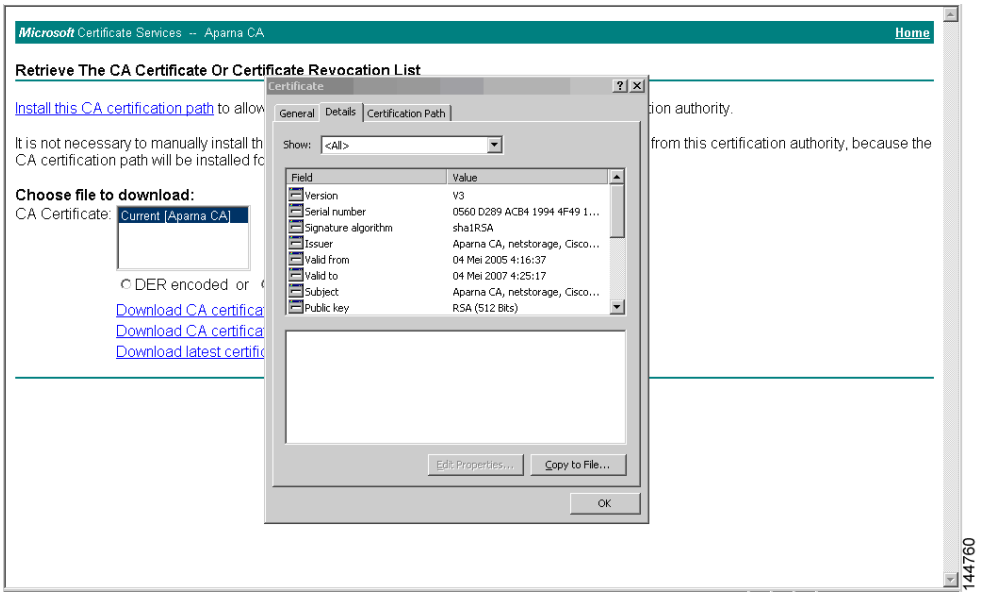

Step 5 Select the **Base-64 encoded X.509 (CER)** on the Certificate Export Wizard dialog box and click **Next**.

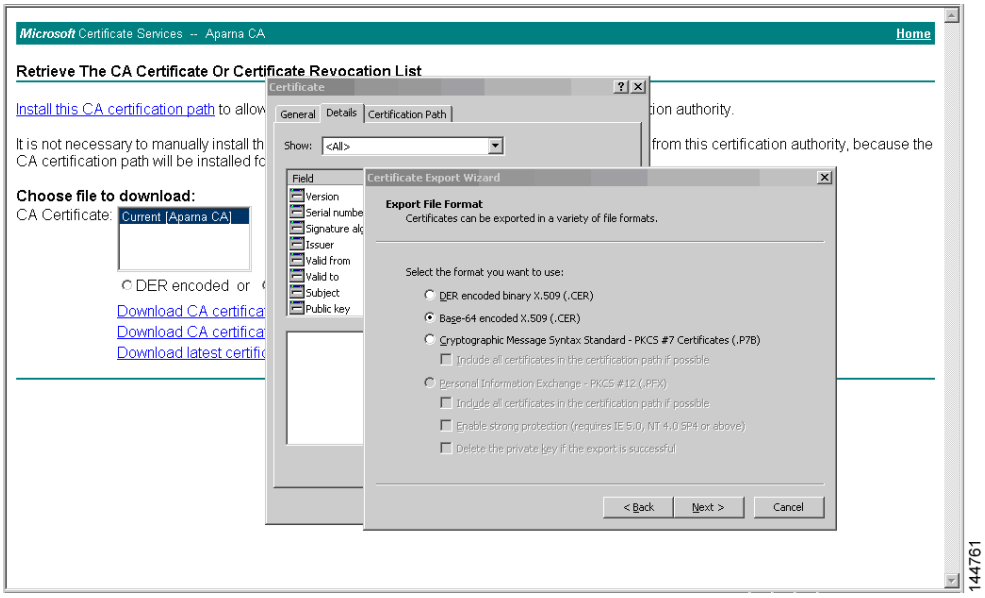

Step 6 Click the **Finish** button on the Certificate Export Wizard dialog box.

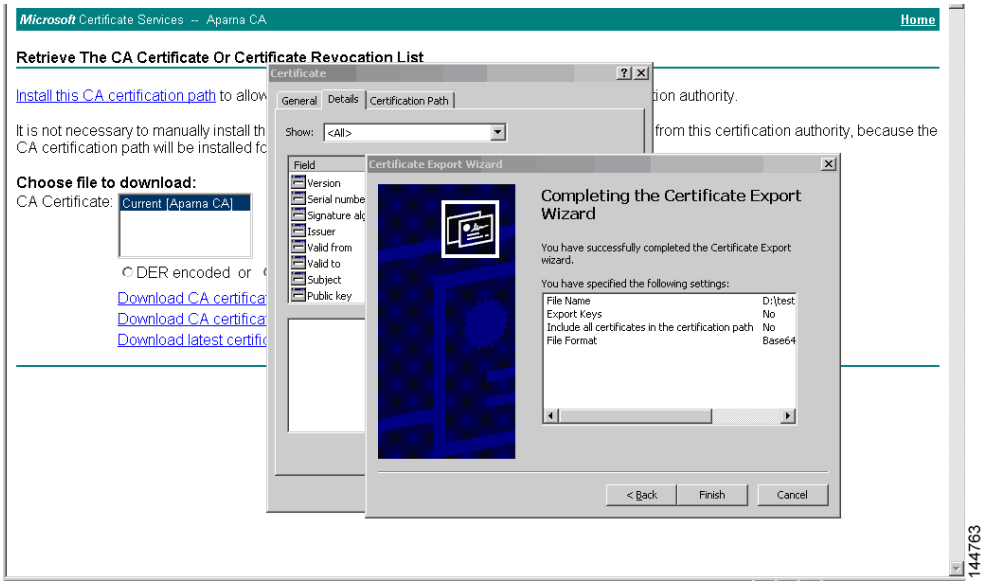

Step 7 Display the CA certificate stored in Base-64 (PEM) format using the Microsoft Windows **type** command.

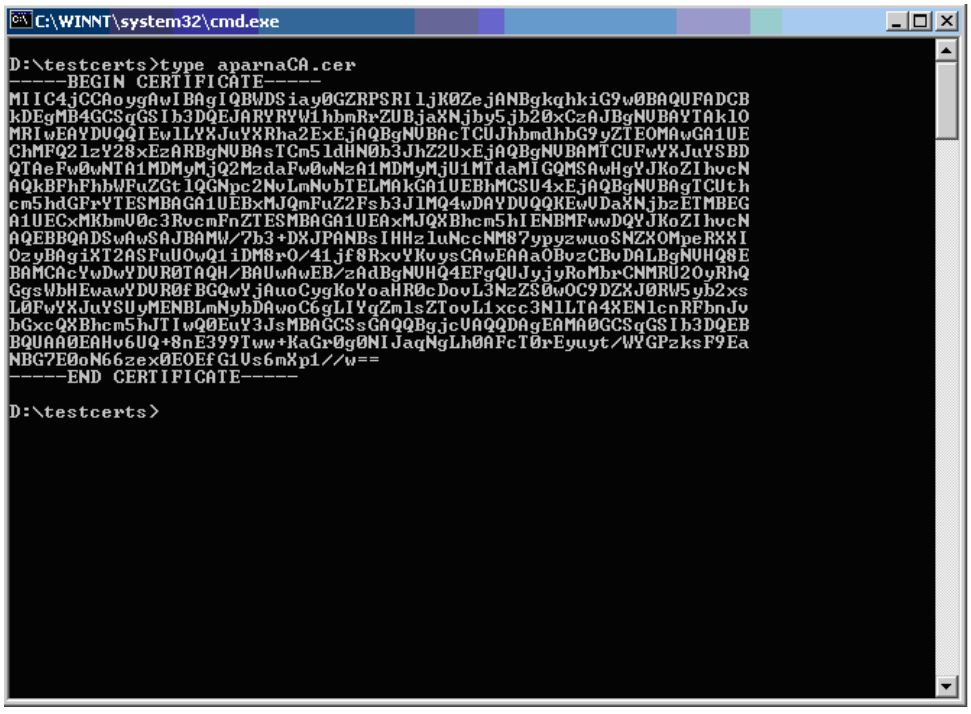

Requesting an Identity Certificate

To request an identify certificate from a Microsoft Certificate server using a PKCS#10 certificate signing request (CRS), follow these steps:

Step 1 Select the Request an identity certificate radio button on the Microsoft Certificate Services web interface and click **Next**.

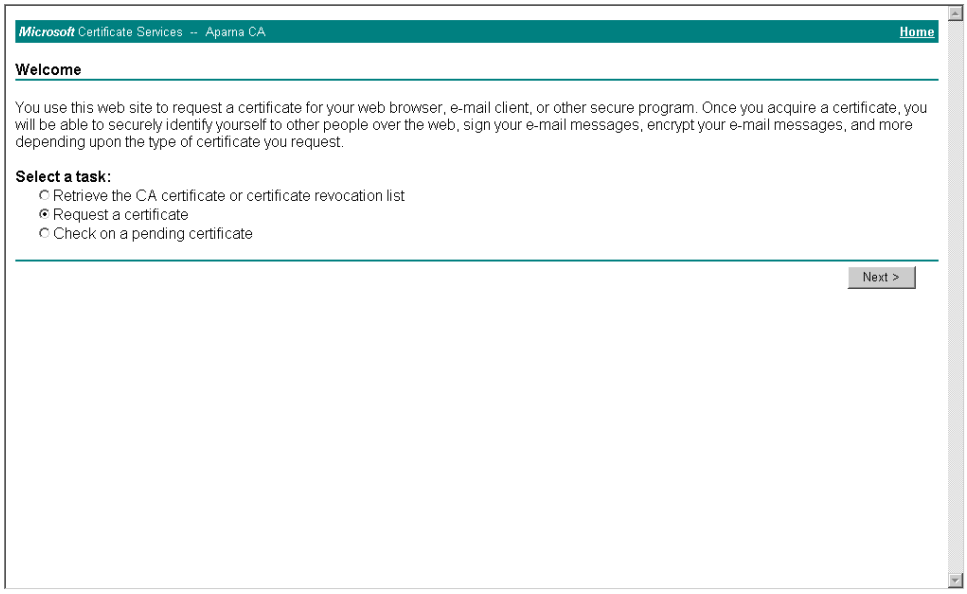

Step 2 Select the **Advanced Request** radio button and click **Next**.

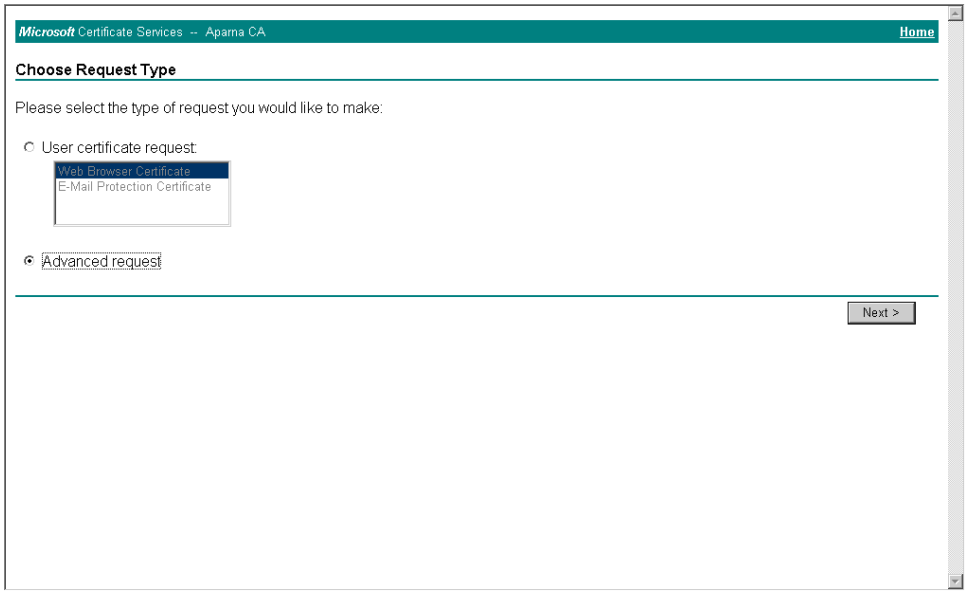

Step 3 Select the **Submit a certificate request using a base64 encoded PKCS#10 file or a renewal request using a base64 encoded PKCS#7 file** radio button and click **Next**.

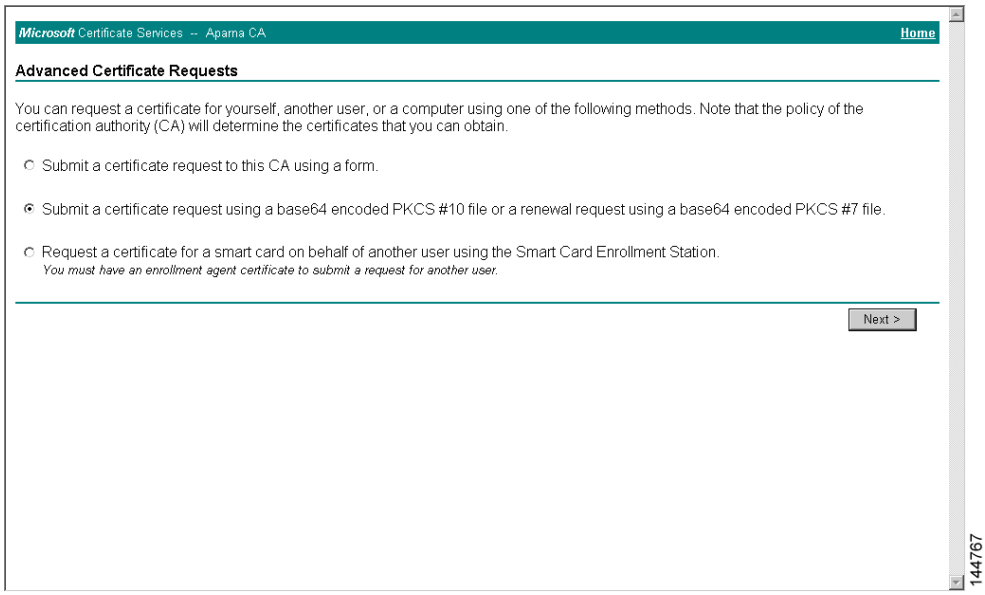

Step 4 Paste the base64 PKCS#10 certificate request in the Saved Request text box and click **Next**. The certificate request is copied from the MDS switch console (see the "Generating Certificate Requests" section on page 35-10 and "Configuring Certificates on the MDS Switch" section on page 35-16)

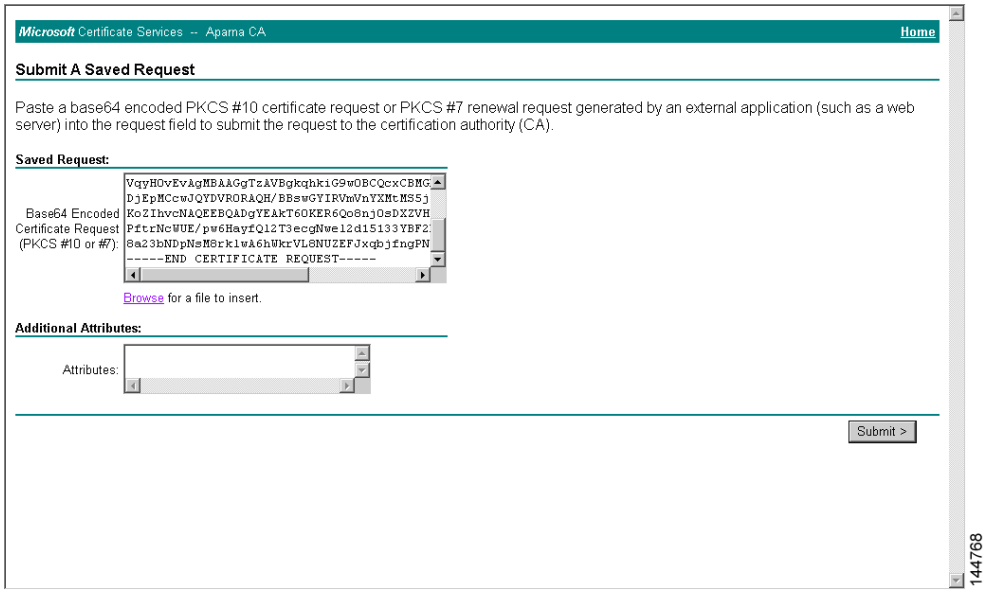

Step 5 Wait one or two days until the certificate is issued by the CA administrator.

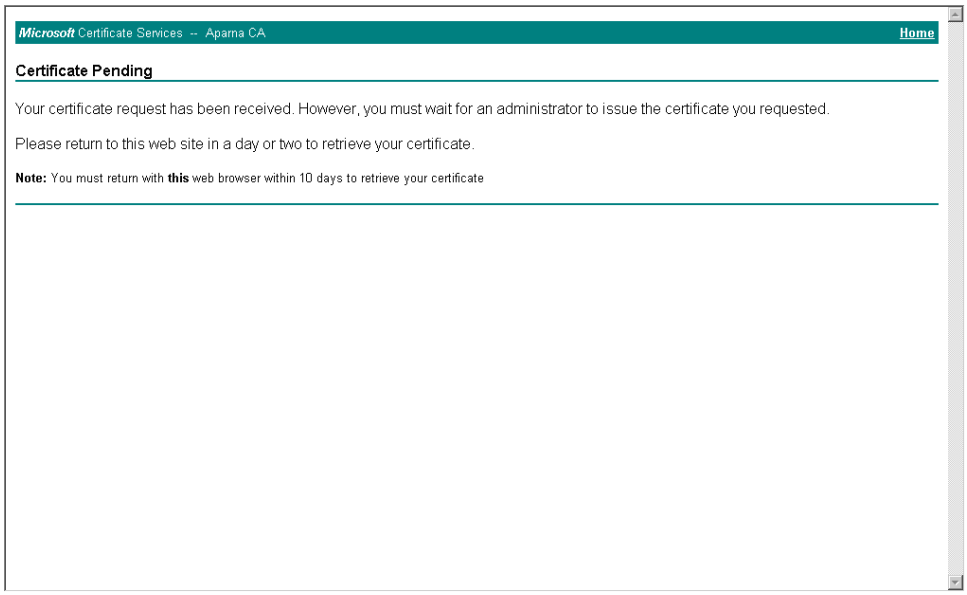

Step 6 The CA administrator approves the certificate request.

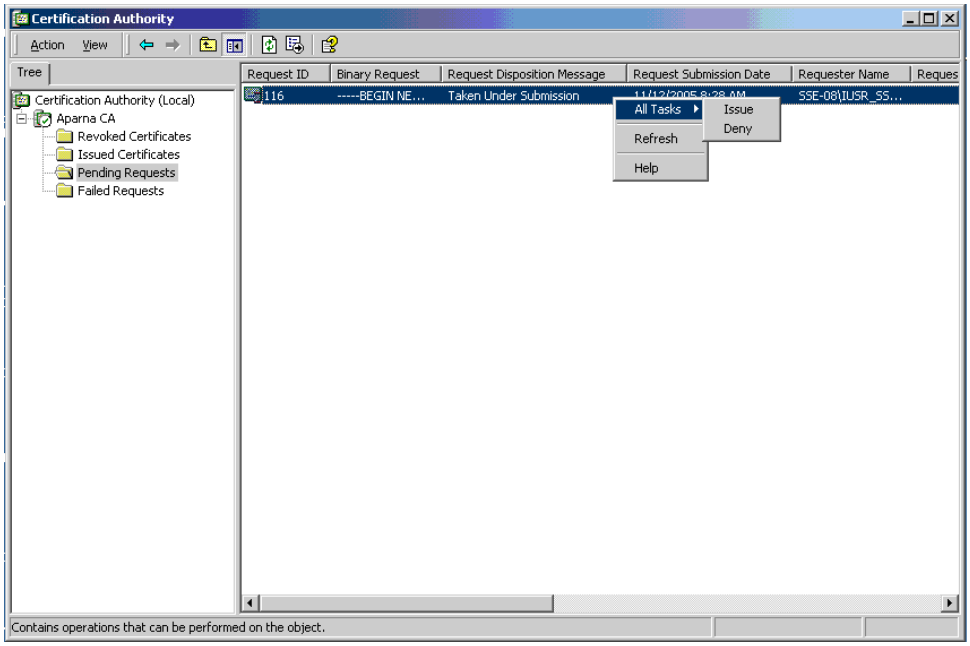

Step 7 Select the **Check on a pending certificate** radio button on the Microsoft Certificate Services web interface and click **Next**.

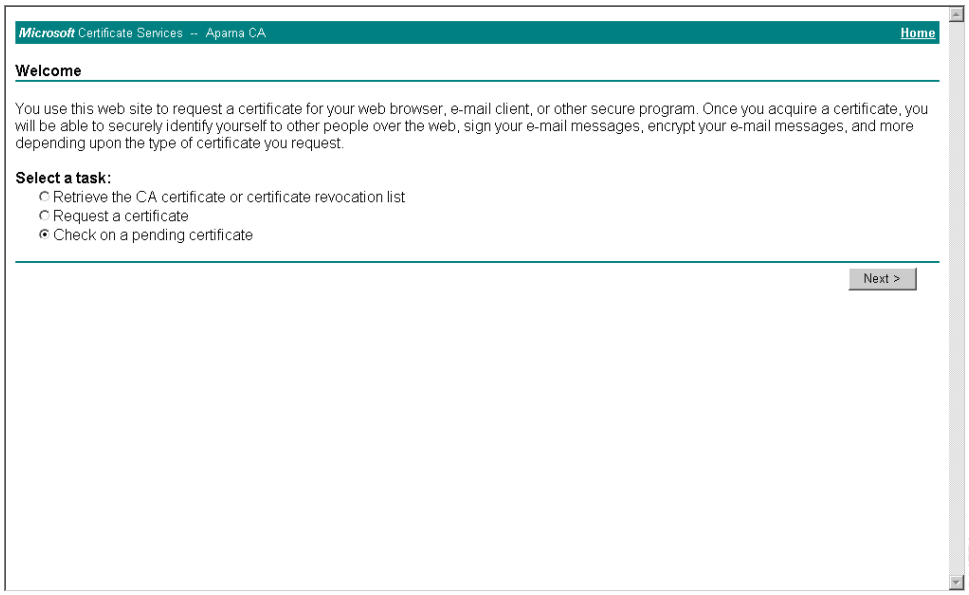

Step 8 Select the certificate request you want to check and click **Next**.

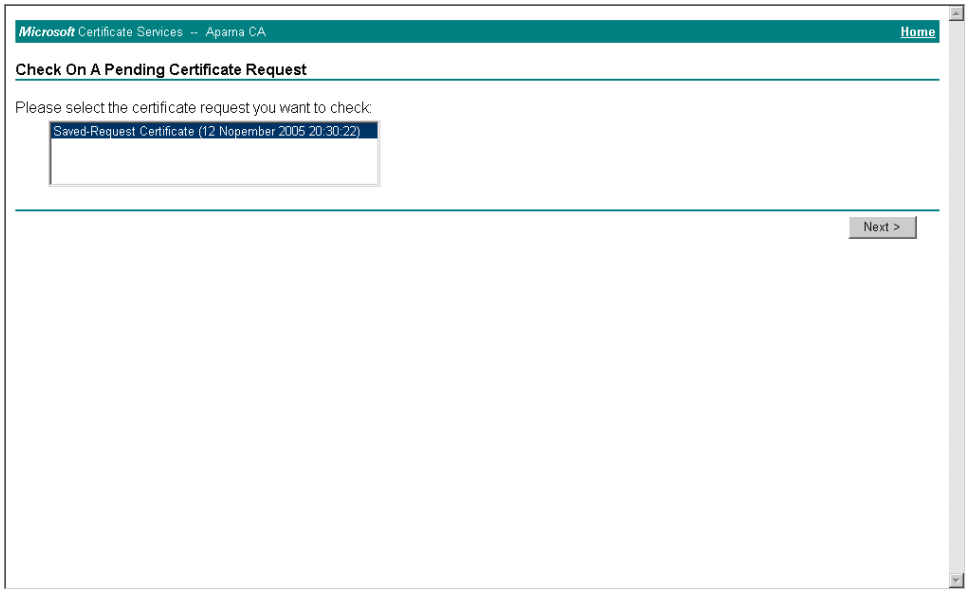

Step 9 Select **Base 64 encoded** and click the **Download CA certificate** link.

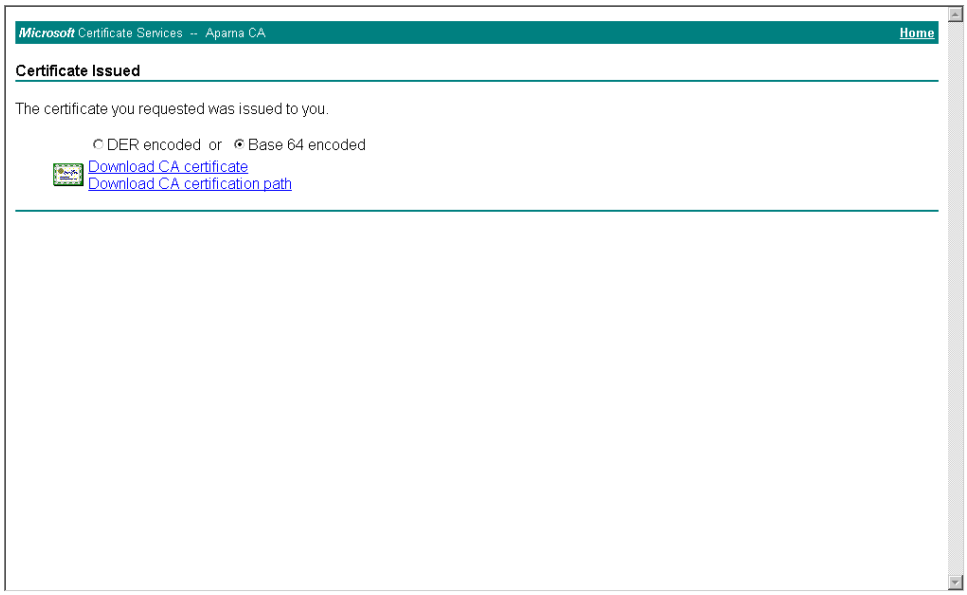

Step 10 Click **Open** on the File Download dialog box.

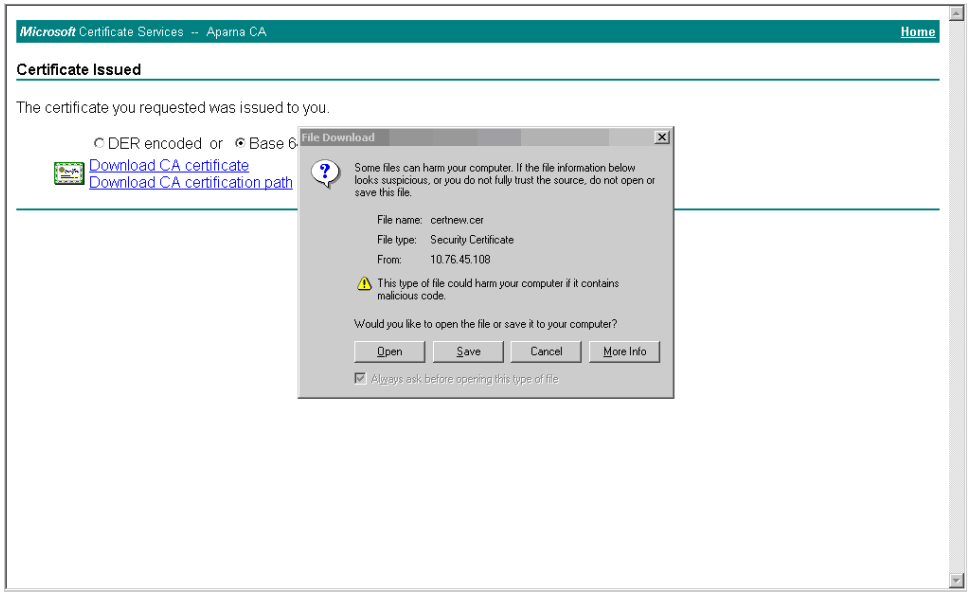

Step 11 Click the **Details** tab on the Certificate dialog and click the **Copy to File** button. Select the **Base-64 encoded X.509 (.CER)** radio button on the Certificate Export Wizard dialog box and click **Next**.

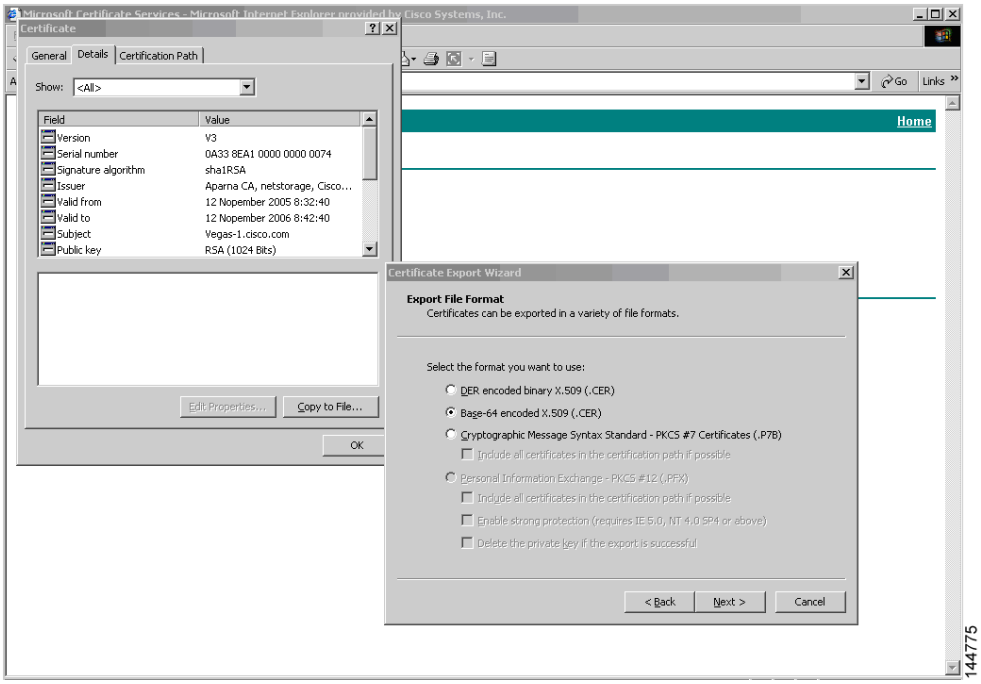

Step 12 Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box, then click **Next**.

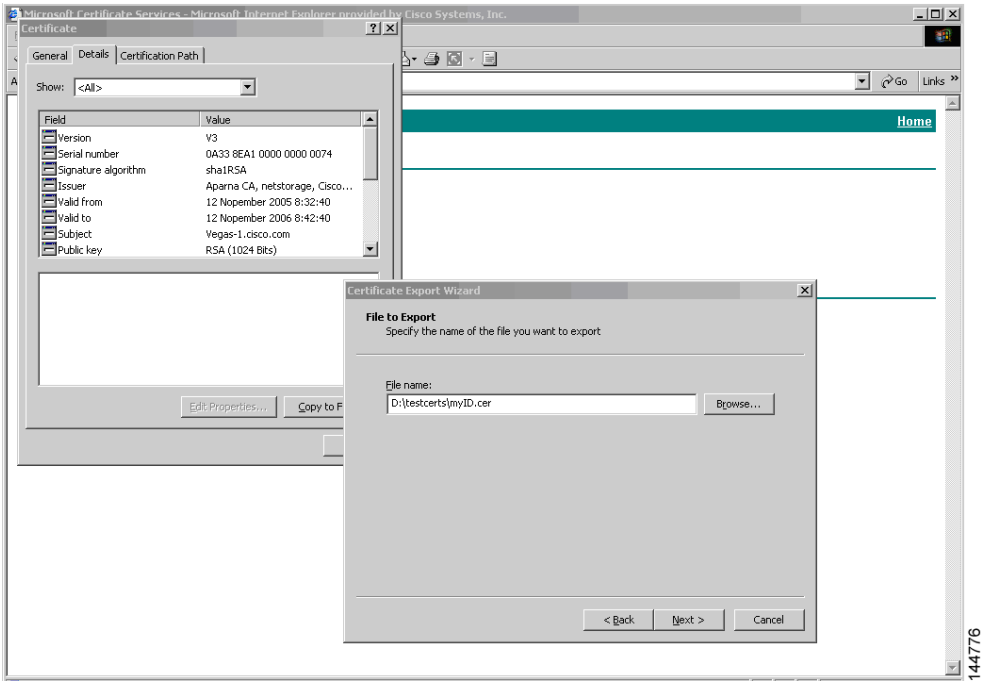

Step 13 Click **Finish**.

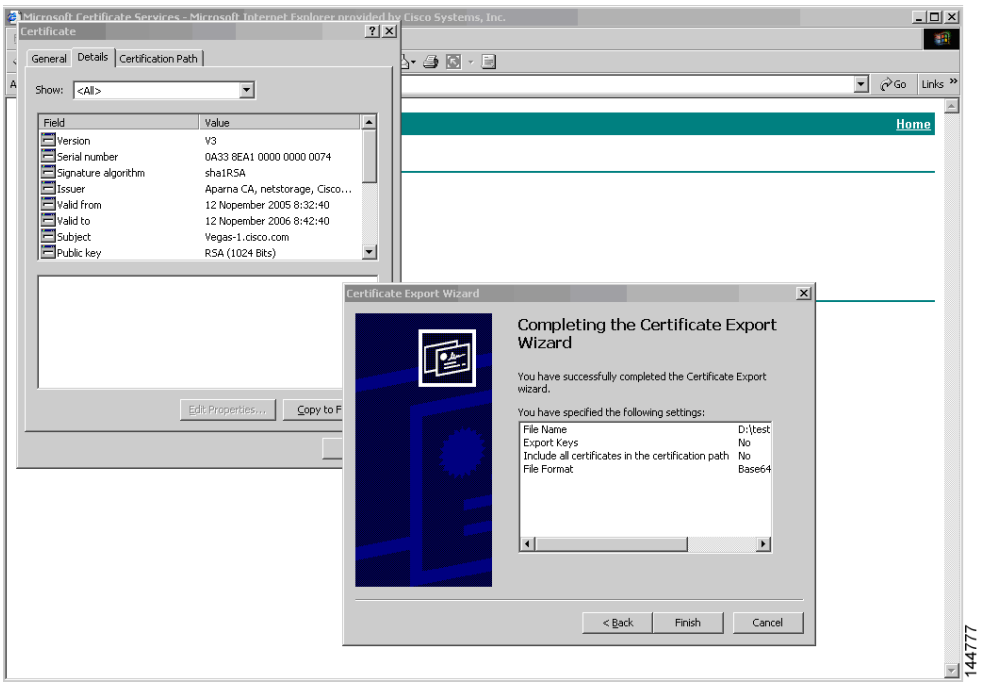

Step 14 Display the identity certificate in base64-encoded format using the Microsoft Windows **type** command.

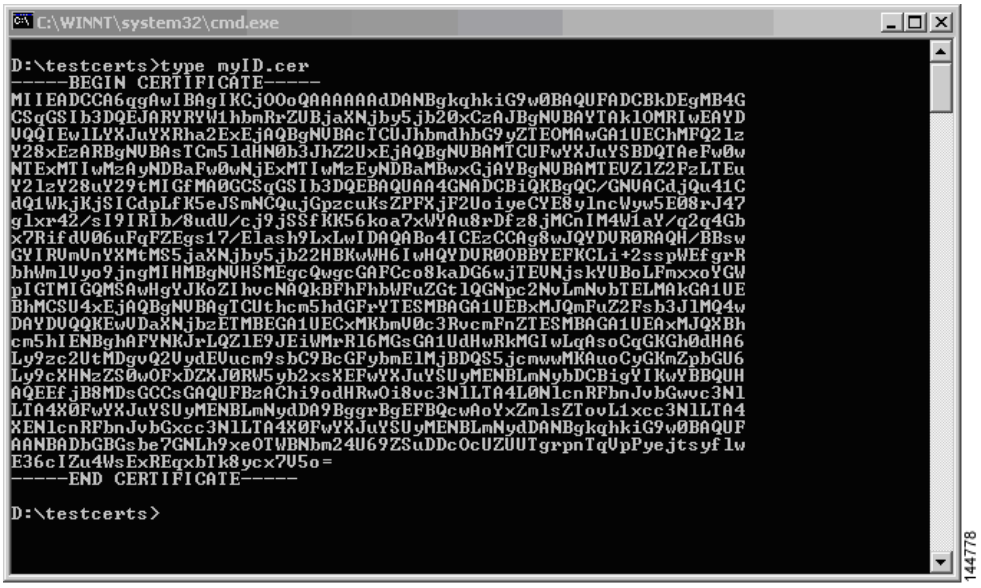

Revoking a Certificate

To revoke a certificate using the Microsoft CA administrator program, follow these steps:

Step 1 Click the **Issued Certificates** folder on the Certification Authority tree. From the list, right-click the certificate you want to revoke.

Step 2 Select **All Tasks > Revoke Certificate**.

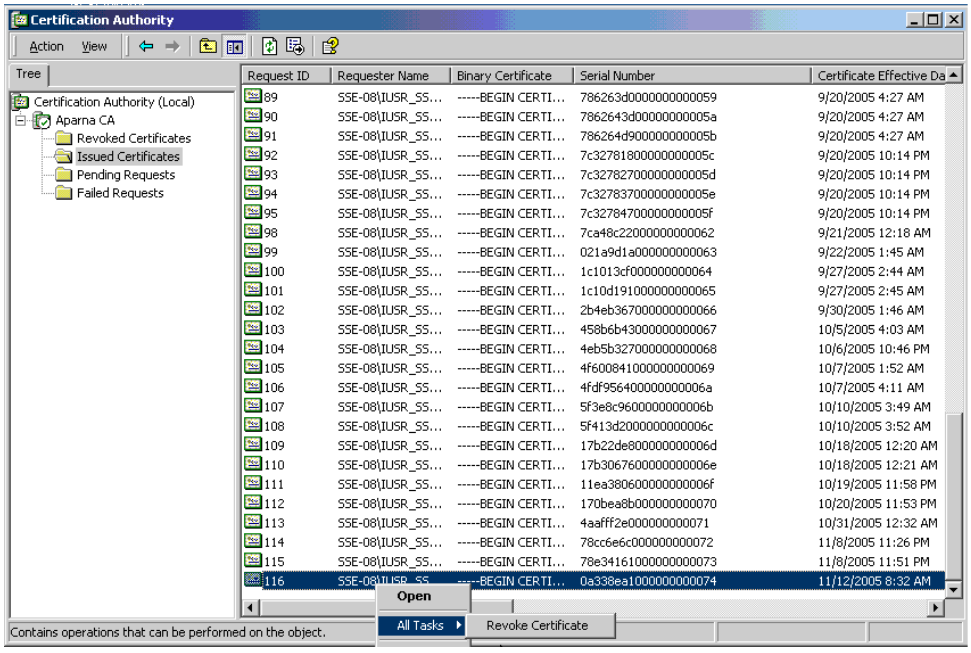

Step 3 Select a reason for the revocation from the Reason code drop-down list, and click **Yes**.

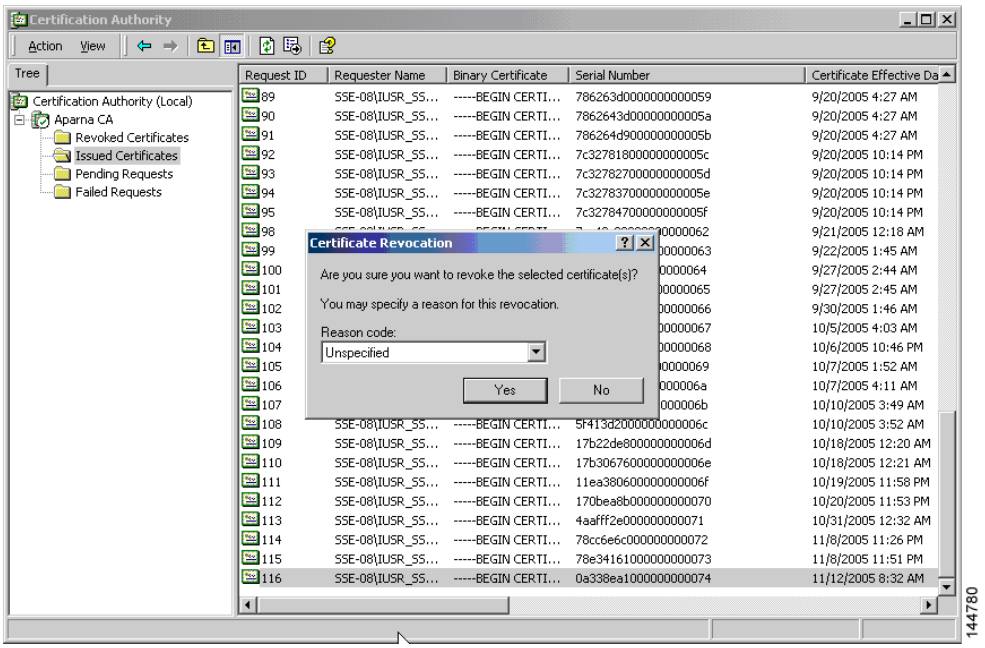

Step 4 Click the **Revoked Certificates** folder to list and verify the certificate revocation.

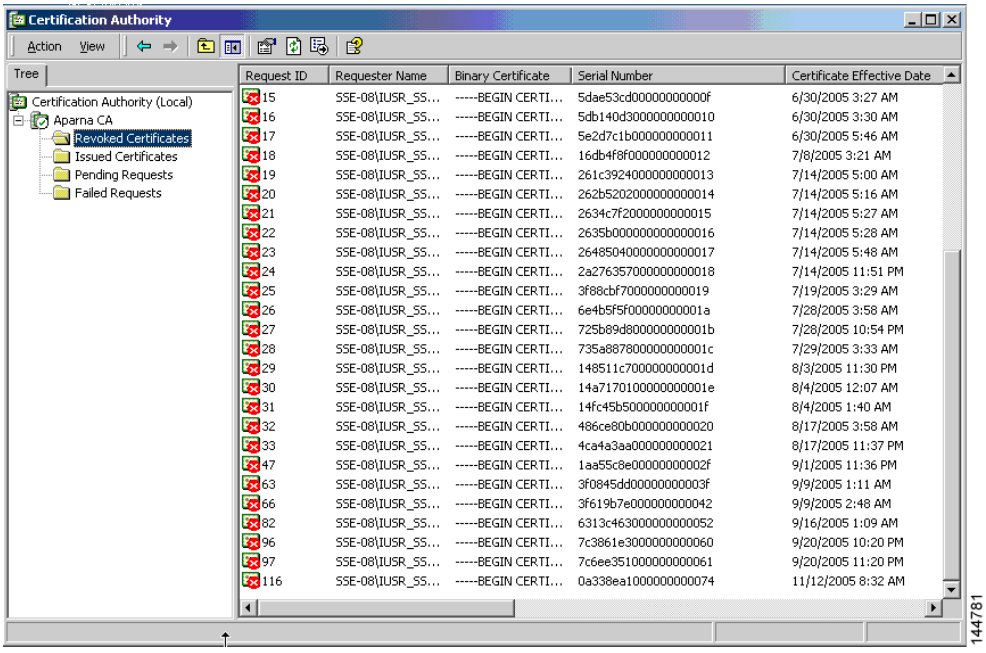

Generating and Publishing the CRL

To generate and publish the CRL using the Microsoft CA administrator program, follow these steps:

Step 1 Select **Action > All Tasks > Publish** on the Certification Authority screen.

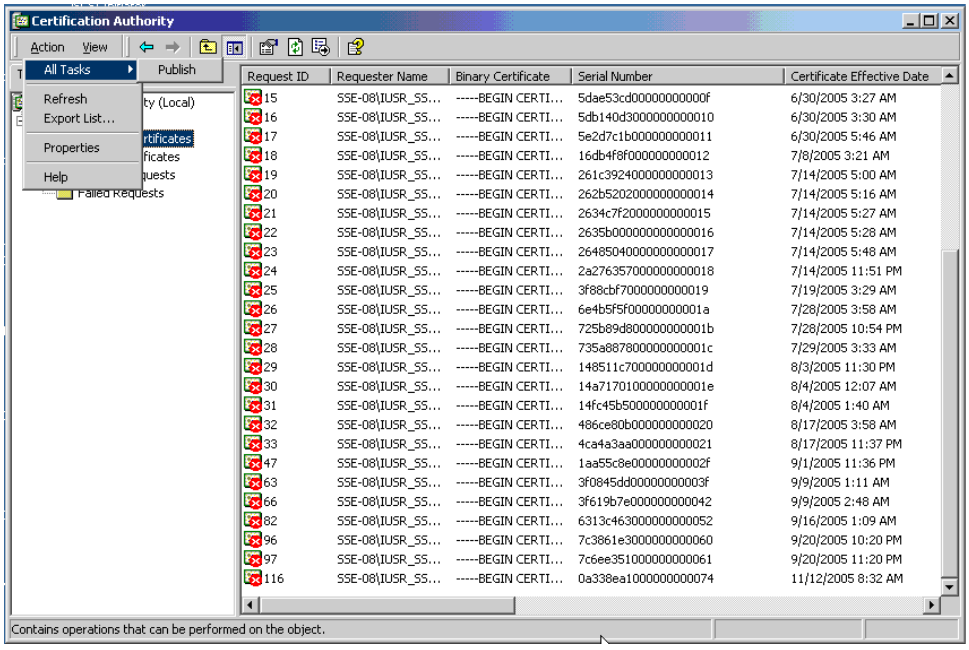

Step 2 Click **Yes** on the Certificate Revocation List dialog box to publish the latest CRL.

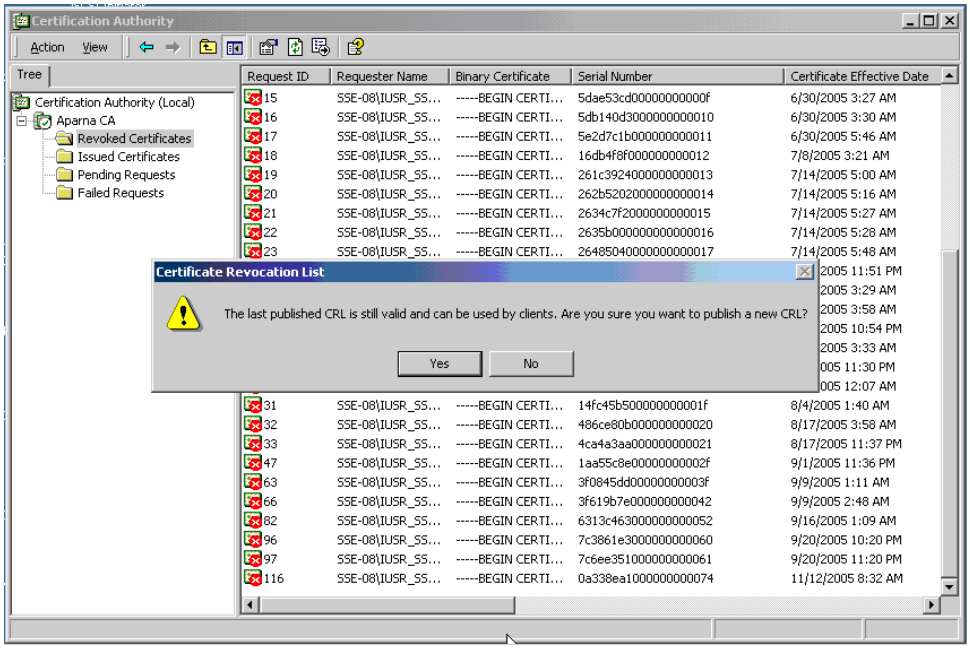

Downloading the CRL

To download the CRL from the Microsoft CA website, follow these steps:.

Step 1 Select **Request the CA certificate or certificate revocation list** radio button on the Microsoft Certificate Services web interface and click **Next**.

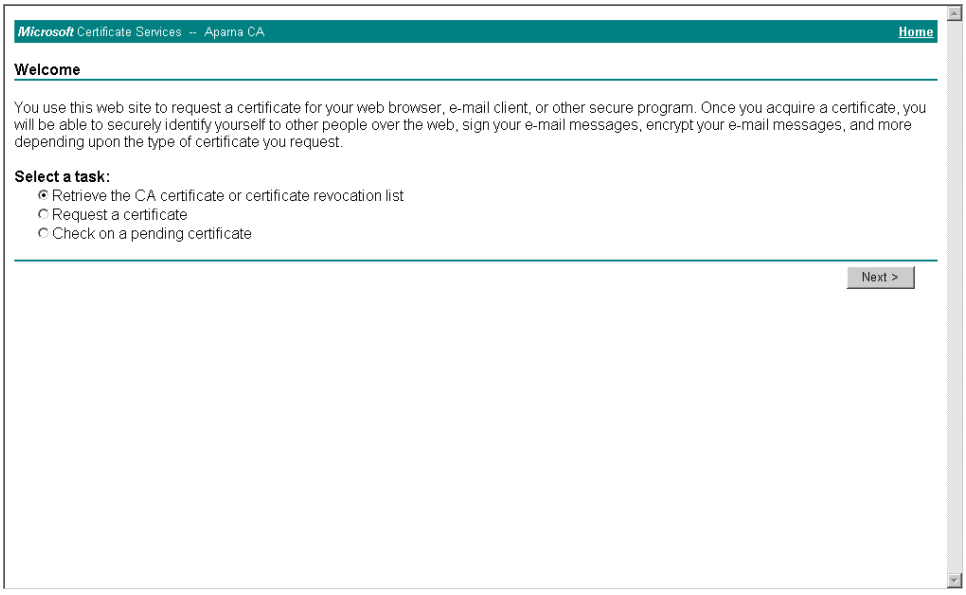

Step 2 Click the **Download latest certificate revocation list** link.

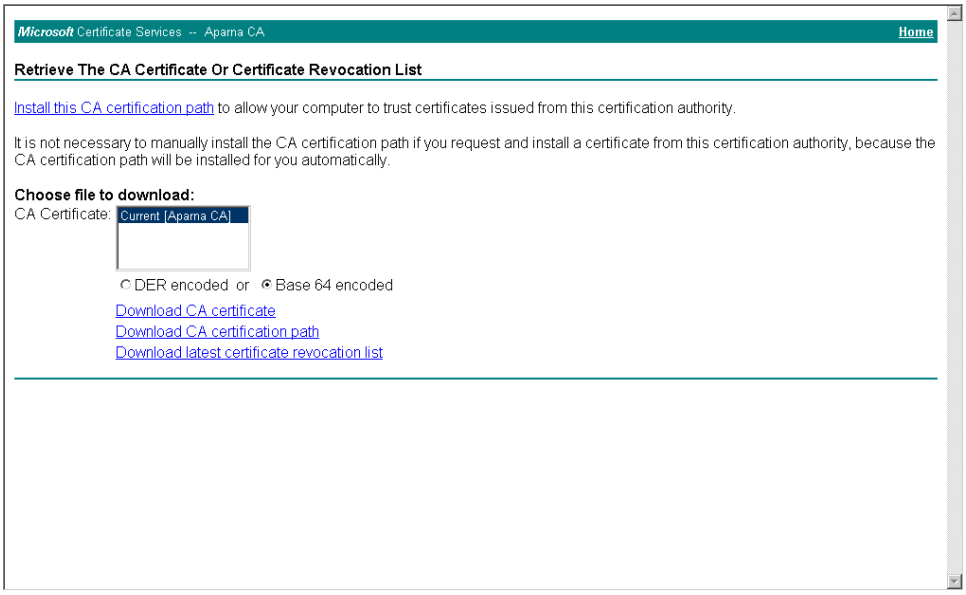

Step 3 Click **Save** in the File Download dialog box.

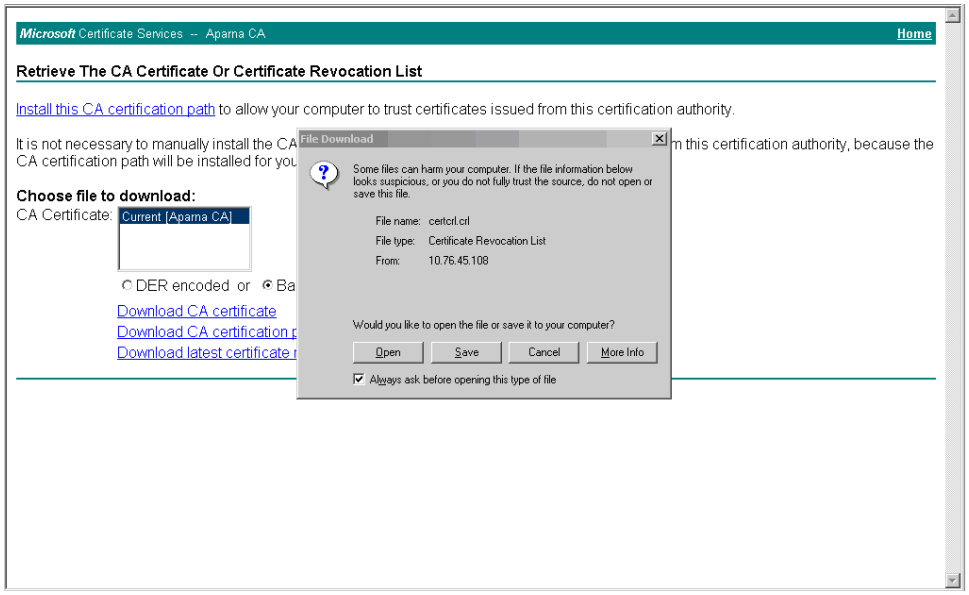

Step 4 Enter the destination file name in the Save As dialog box and click **Save**.

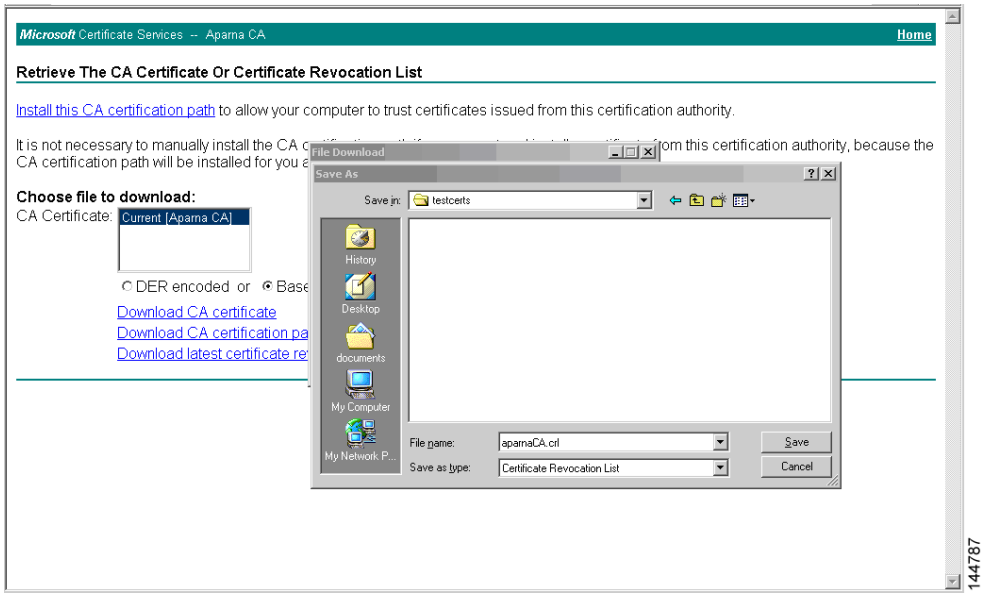

Step 5 Display the CRL using the Microsoft Windows **type** command.

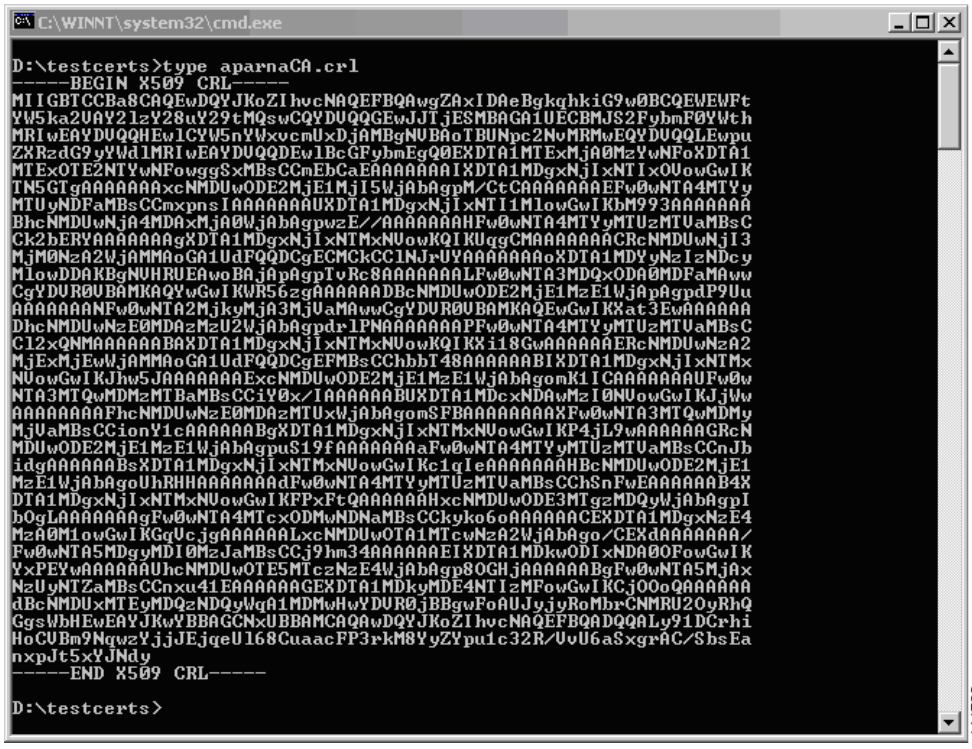

Importing the CRL

To import the CRL to the trust point corresponding to the CA, follow these steps:

Step 1 Copy the CRL file to the MDS switch bootflash.

```
Vegas-1# copy tftp:apranaCA.crl bootflash:aparnaCA.crl
```

Step 2 Configure the CRL.

```
Vegas-1# config t
Vegas-1(config)# crypto ca crl request myCA bootflash:aparnaCA.crl
Vegas-1(config)#
```

Step 3 Display the contents of the CRL.

```
Vegas-1(config)# do sh crypto ca crl myCA
Trustpoint: myCA
CRL:
Certificate Revocation List (CRL):
        Version 2 (0x1)
        Signature Algorithm: sha1WithRSAEncryption
        Issuer: /emailAddress=admin@yourcompany.com/C=IN/ST=Karnatak
Yourcompany/OU=netstorage/CN=Aparna CA
        Last Update: Nov 12 04:36:04 2005 GMT
        Next Update: Nov 19 16:56:04 2005 GMT
        CRL extensions:
            X509v3 Authority Key Identifier:
            keyid:27:28:F2:46:83:1B:AC:23:4C:45:4D:8E:C9:18:50:1

            1.3.6.1.4.1.311.21.1:
                ...
Revoked Certificates:
    Serial Number: 611B09A1000000000002
        Revocation Date: Aug 16 21:52:19 2005 GMT
Serial Number: 4CDE464E000000000003
        Revocation Date: Aug 16 21:52:29 2005 GMT
    Serial Number: 4CFC2B42000000000004
        Revocation Date: Aug 16 21:52:41 2005 GMT
    Serial Number: 6C699EC2000000000005
        Revocation Date: Aug 16 21:52:52 2005 GMT
    Serial Number: 6CCF7DDC000000000006
        Revocation Date: Jun  8 00:12:04 2005 GMT
    Serial Number: 70CC4FFF000000000007
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 4D9B1116000000000008
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 52A80230000000000009
        Revocation Date: Jun 27 23:47:06 2005 GMT
        CRL entry extensions:
            X509v3 CRL Reason Code:
            CA Compromise
Serial Number: 5349AD4600000000000A
        Revocation Date: Jun 27 23:47:22 2005 GMT
        CRL entry extensions:
            X509v3 CRL Reason Code:
            CA Compromise
Serial Number: 53BD173C00000000000B
        Revocation Date: Jul  4 18:04:01 2005 GMT
        CRL entry extensions:
            X509v3 CRL Reason Code:
            Certificate Hold
Serial Number: 591E7ACE00000000000C
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 5D3FD52E00000000000D
        Revocation Date: Jun 29 22:07:25 2005 GMT
        CRL entry extensions:
            X509v3 CRL Reason Code:
            Key Compromise
Serial Number: 5DAB771300000000000E
        Revocation Date: Jul 14 00:33:56 2005 GMT
    Serial Number: 5DAE53CD00000000000F
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 5DB140D3000000000010
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 5E2D7C1B000000000011
        Revocation Date: Jul  6 21:12:10 2005 GMT
        CRL entry extensions:
```

```
            X509v3 CRL Reason Code:
            Cessation Of Operation
Serial Number: 16DB4F8F000000000012
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 261C3924000000000013
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 262B5202000000000014
        Revocation Date: Jul 14 00:33:10 2005 GMT
    Serial Number: 2634C7F2000000000015
        Revocation Date: Jul 14 00:32:45 2005 GMT
    Serial Number: 2635B000000000000016
        Revocation Date: Jul 14 00:31:51 2005 GMT
    Serial Number: 26485040000000000017
        Revocation Date: Jul 14 00:32:25 2005 GMT
    Serial Number: 2A276357000000000018
Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 3F88CBF7000000000019
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 6E4B5F5F00000000001A
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 725B89D800000000001B
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 735A887800000000001C
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 148511C700000000001D
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 14A7170100000000001E
        Revocation Date: Aug 16 21:53:15 2005 GMT
    Serial Number: 14FC45B500000000001F
        Revocation Date: Aug 17 18:30:42 2005 GMT
    Serial Number: 486CE80B000000000020
        Revocation Date: Aug 17 18:30:43 2005 GMT
    Serial Number: 4CA4A3AA000000000021
        Revocation Date: Aug 17 18:30:43 2005 GMT
    Serial Number: 1AA55C8E00000000002F
        Revocation Date: Sep  5 17:07:06 2005 GMT
    Serial Number: 3F0845DD00000000003F
        Revocation Date: Sep  8 20:24:32 2005 GMT
    Serial Number: 3F619B7E000000000042
        Revocation Date: Sep  8 21:40:48 2005 GMT
    Serial Number: 6313C463000000000052
        Revocation Date: Sep 19 17:37:18 2005 GMT
Serial Number: 7C3861E3000000000060
        Revocation Date: Sep 20 17:52:56 2005 GMT
    Serial Number: 7C6EE351000000000061
        Revocation Date: Sep 20 18:52:30 2005 GMT
    Serial Number: 0A338EA1000000000074       <-- Revoked identity certificate
        Revocation Date: Nov 12 04:34:42 2005 GMT
    Signature Algorithm: sha1WithRSAEncryption
        0b:cb:dd:43:0a:b8:62:1e:80:95:06:6f:4d:ab:0c:d8:8e:32:
        44:8e:a7:94:97:af:02:b9:a6:9c:14:fd:eb:90:cf:18:c9:96:
        29:bb:57:37:d9:1f:d5:bd:4e:9a:4b:18:2b:00:2f:d2:6e:c1:
        1a:9f:1a:49:b7:9c:58:24:d7:72
```

Note The identity certificate for the switch that was revoked (serial number 0A338EA1000000000074) is listed at the end.

Maximum Limits

Table 35-1 lists the maximum limits for CAs and digital certificate parameters.

Table 35-1 Maximum Limits for CA and Digital Certificate

Feature	Maximum Limit
Trust points declared on a switch	16.
RSA key-pairs generated on a switch	16.
Identity certificates configured on a switch	16.
Certificates in a CA certificate chain	10.
Trust points authenticated to a specific CA	10.

Default Settings

Table 35-2 lists the default settings for CAs and digital certificate parameters.

Table 35-2 Default CA and Digital Certificate Parameters

Parameters	Default
Trust point	None.
RSA key-pair	None.
RSA key-pair label	Switch FQDN.
RSA key-pair modulus	512.
RSA key-pair exportable	Yes.
Revocation check method of trust point	CRL.

CHAPTER 36

Configuring IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is the latest version of RFC 2401. Cisco SAN-OS IPsec implements RFC 2402 through RFC 2410.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

Note The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is other times used to describe only the data services.

This chapter includes the following sections:

About IPsec

Note IPsec is not supported by the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec devices:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.
- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.

Note The term *data authentication* is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco SAN-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.

Note The Encapsulating Security Payload (ESP) protocol is a header inserted into an existing TCP/IP packet, the size of which depends on the actual encryption and authentication algorithms negotiated. To avoid fragmentation, the encrypted packet fits into the interface maximum transmission unit (MTU). The path MTU calculation for TCP takes into account the addition of ESP headers, plus the outer IP header in tunnel mode, for encryption. The MDS switches allow 100 bytes for packet growth for IPsec encryption.

Figure 36-1 shows different IPsec scenarios.

Figure 36-1 FCIP and iSCSI Scenarios Using MPS-14/2 Modules

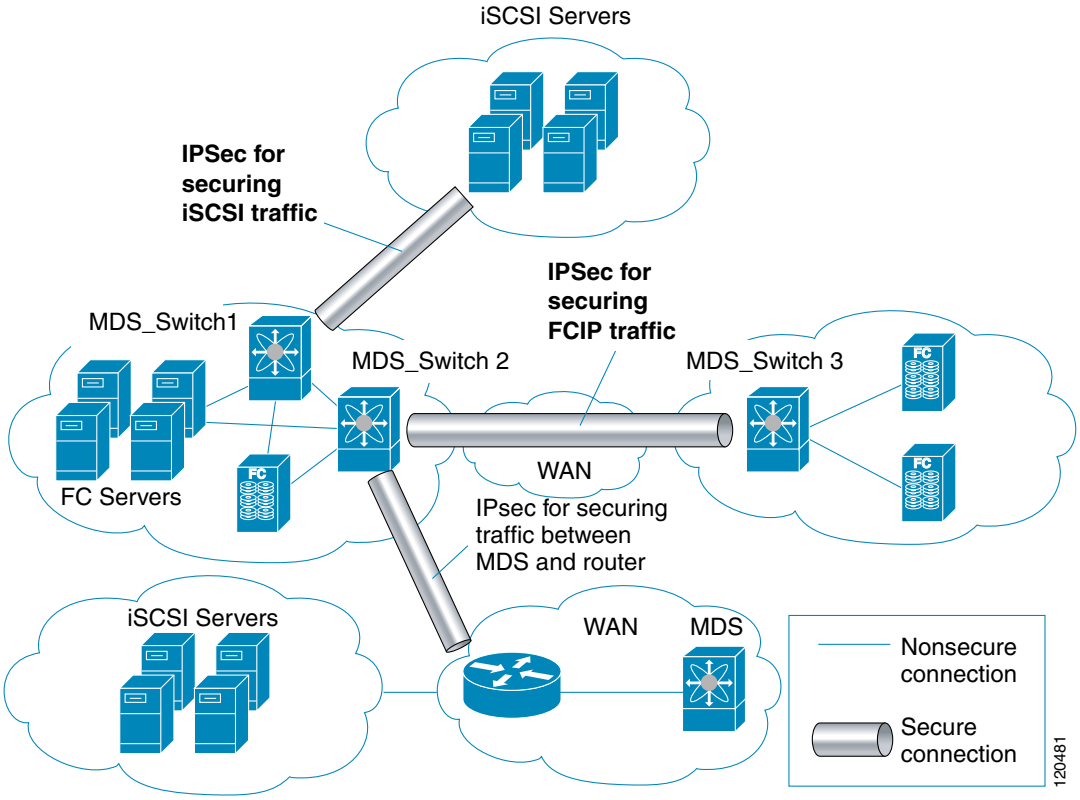

About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

Note IKE is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

IPsec Prerequisites

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE_PKG license (see Chapter 3, "Obtaining and Installing Licenses").
- Configure IKE as described in the "About IKE Initialization" section on page 36-11.

Note The IPsec feature inserts new headers in existing packets (see the"Configuring the MTU Frame Size" section on page 45-3 for more information).

Using IPsec

To use the IPsec feature, follow these steps:

Step 1 Obtain the ENTERPRISE_PKG license to enable IPSEC for iSCSI to enable IPsec for FCIP. See Chapter 3, "Obtaining and Installing Licenses."

Step 2 Configure IKE as described in the "Manually Configuring IPsec and IKE" section on page 36-10.

Note The IPsec feature inserts new headers in existing packets (see the "Configuring the MTU Frame Size" section on page 45-3 for more information).

This section contains the following topics:

IPsec Compatibility

IPsec features are compatible with the following Cisco MDS 9000 Family hardware:

- Cisco 14/2-port Multiprotocol Services (MPS-14/2) modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.
- The IPsec feature is not supported on the management interface.

IPsec features are compatible with the following fabric setup:

- Two connected Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later.
- A Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later connected to any IPsec compliant device.
- The following features are not supported in the Cisco SAN-OS implementation of the IPsec feature:
  - Authentication Header (AH).
  - Transport mode.
  - Security association bundling.
  - Manually configuring security associations.
  - Per host security association option in a crypto map.
  - Security association idle timeout
  - Dynamic crypto maps.

Note Any reference to crypto maps in this document, only refers to static crypto maps.

IPsec and IKE Terminology

The terms used in this chapter are explained in this section.

- Security association (SA)— An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
  - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
  - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
  - Transform—A list of operations done to provide data authentication and data confidentiality. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
  - Session key—The key used by the transform to provide security services.
  - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and, if required, is automatically renegotiated (rekeyed).
  - Mode of operation—Two modes of operation are generally available for IPsec: tunnel mode and transport mode. The Cisco SAN-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet, including its header. The gateways encrypt traffic on behalf of the hosts and subnets.
    The Cisco SAN-OS implementation of IPsec does not support transport mode.

Note The term *tunnel mode* is different from the term *tunnel*, which is used to indicate a secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.
- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
  - Data integrity—Verifies that data has not been altered.
  - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address and mask or prefix, destination address mask or prefix length, IP next protocol field, and source and destination ports, where the protocol and port fields can have any of these values. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.
- Perfect forward secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—An ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if it should be allowed in clear text, or if it should be dropped.
  - The IPsec SPDs are derived from user configuration of crypto maps.
  - The IKE SPD is configured by the user.

Supported IPsec Transforms and Algorithms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Note Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.

- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm.

Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol that allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Note Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- The switch authentication algorithm uses the preshared keys based on the IP address (see "Setting Transmission Retry Count for the RADIUS Server" section on page 33-11 for more information on preshared keys).

IPsec Digital Certificate Support

This section describes the advantages of using certificate authorities (CAs) and digital certificates for authentication.

For more information on CAs and digital certificates, see Chapter 35, "Configuring Certificate Authorities and Digital Certificates."

Implementing IPsec Without CAs and Digital Certificates

Without a CA and digital certificates, enabling IPsec services (such as encryption) between two Cisco MDS switches requires that each switch has the key of the other switch (such as an RSA public key or a shared key). You must manually specify either the RSA public keys or preshared keys on each switch in the fabric using IPsec services. Also, each new device added to the fabric will require manual configuration of the other switches in the fabric to support secure communication.

In Figure 36-2, each switch uses the key of the other switch to authenticate the identity of the other switch; this authentication always occurs when IPsec traffic is exchanged between the two switches.

If you have multiple Cisco MDS switches in a mesh topology and wish to exchange IPsec traffic passing among all of those switches, you must first configure shared keys or RSA public keys among all of those switches.

Figure 36-2 Two IPsec Switches Without CAs and Digital Certificates

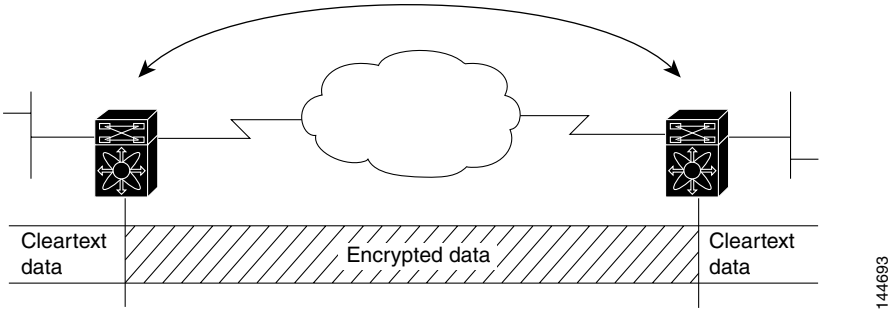

Every time a new switch is added to the IPsec network, you must configure keys between the new switch and each of the existing switches. (In Figure 36-3, four additional two-part key configurations are required to add a single encrypting switch to the network.)

Consequently, the more devices that require IPsec services, the more involved the key administration becomes. This approach does not scale well for larger, more complex encrypting networks.

Figure 36-3 Four IPsec Switches Without a CA and Digital Certificates

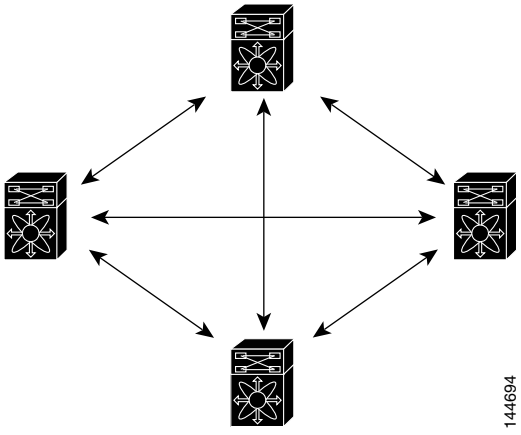

Implementing IPsec with CAs and Digital Certificates

With CA and digital certificates, you do not have to configure keys between all the encrypting switches. Instead, you individually enroll each participating switch with the CA, requesting a certificate for the switch. When this has been accomplished, each participating switch can dynamically authenticate all the other participating switches. When two devices want to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new device is added to the network, you simply enroll that device with a CA, and none of the other devices needs modification. When the new device attempts an IPsec connection, certificates are automatically exchanged and the device can be authenticated.

Figure 36-4 shows the process of dynamically authenticating the devices.

Figure 36-4 Dynamically Authenticating Devices with a CA

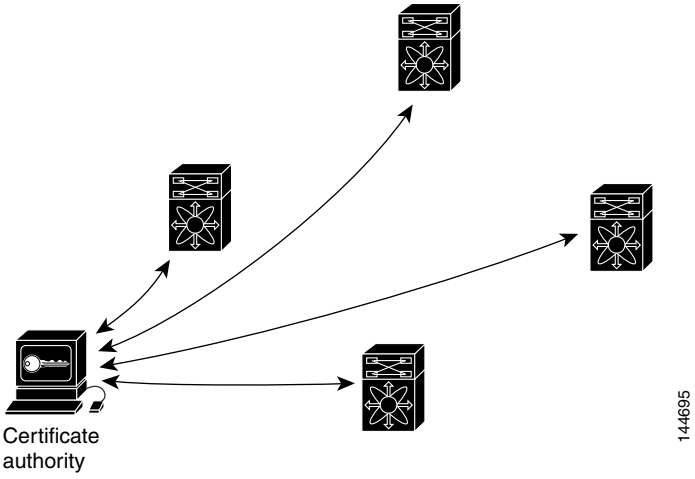

To add a new IPsec switch to the network, you need only configure that new switch to request a certificate from the CA, instead of making multiple key configurations with all the other existing IPsec switches.

How CA Certificates Are Used by IPsec Devices

When two IPsec switches want to exchange IPsec-protected traffic passing between them, they must first authenticate each other—otherwise, IPsec protection cannot occur. The authentication is done with IKE.

IKE can use two methods to authenticate the switches, using preshared keys without a CA and using RSA key-pairs with a CA. Both methods require that keys must be preconfigured between the two switches.

Without a CA, a switch authenticates itself to the remote switch using either RSA-encrypted preshared keys.

With a CA, a switch authenticates itself to the remote switch by sending a certificate to the remote switch and performing some public key cryptography. Each switch must send its own unique certificate that was issued and validated by the CA. This process works because the certificate of each switch encapsulates the public key of the switch, each certificate is authenticated by the CA, and all participating switches recognize the CA as an authenticating authority. This scheme is called IKE with an RSA signature.

Your switch can continue sending its own certificate for multiple IPsec sessions, and to multiple IPsec peers until the certificate expires. When the certificate expires, the switch administrator must obtain a new one from the CA.

CAs can also revoke certificates for devices that will no longer participate in IPsec. Revoked certificates are not recognized as valid by other IPsec devices. Revoked certificates are listed in a certificate revocation list (CRL), which each peer may check before accepting a certificate from another peer.

Certificate support for IKE has the following considerations:

- The switch FQDN (host name and domain name) must be configured before installing certificates for IKE.
- Only those certificates that are configured for IKE or general usage are used by IKE.
- The first IKE or general usage certificate configured on the switch is used as the default certificate by IKE.
- The default certificate is for all IKE peers unless the peer specifies another certificate.
- If the peer asks for a certificate which is signed by a CA that it trusts, then IKE uses that certificate, if it exists on the switch, even if it is not the default certificate.
- If the default certificate is deleted, the next IKE or general usage certificate, if any exists, is used by IKE as the default certificate.
- Certificate chaining is not supported by IKE.
- IKE only sends the identity certificate, not the entire CA chain. For the certificate to be verified on the peer, the same CA chain must also exist there.

Manually Configuring IPsec and IKE

This section describes how to manually configure IPsec and IKE .

IPsec provides secure data flows between participating peers. Multiple IPsec data flows can exist between two peers to secure different data flows, with each tunnel using a separate set of SAs.

After you have completed IKE configuration, configure IPsec.

To configure IPsec in each participating IPsec peer, follow these steps:

Step 1 Identify the peers for the traffic to which secure tunnels should be established.

Step 2 Configure the transform set with the required protocols and algorithms.

Step 3 Create the crypto map and apply access control lists (IPv4-ACLs), transform sets, peers, and lifetime values as applicable.

Step 4 Apply the crypto map to the required interface.

This section contains the following topics:

- Configuring an IKE Policy, page 36-13

About IKE Initialization

The IKE feature must first be enabled and configured so the IPsec feature can establish data flow with the required peer. Fabric Manager initializes IKE when you first configure it.

You cannot disable IKE if IPsec is enabled. If you disable the IKE feature, the IKE configuration is cleared from the running configuration.

Enabling IKE

To enable IKE, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike enable`	Enables the IKE feature.
	`switch(config)# no crypto ike enable`	Disables (default) the IKE feature. **Note** You must disable IPsec before you can disable the IKE feature.

About the IKE Domain

You must apply the IKE configuration to an IPsec domain to allow traffic to reach the supervisor module in the local switch. Fabric Manager sets the IPsec domain automatically when you configure IKE.

Configuring the IKE Domain

You must apply the IKE configurations to an IPsec domain to allow traffic to reach the supervisor module in the local switch.

To configure the IPsec domain, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike domain ipsec` `switch(config-ike-ipsec)#`	Allows IKE configurations for IPsec domains.

About IKE Tunnels

An IKE tunnel is a secure IKE session between two endpoints. IKE creates this tunnel to protect IKE messages used in IPsec SA negotiations.

Two versions of IKE are used in the Cisco SAN-OS implementation.

- IKE version 1 (IKEv1) is implemented using RFC 2407, 2408, 2409, and 2412.
- IKE version 2 (IKEv2) is a simplified and more efficient version and does not interoperate with IKEv1. IKEv2 is implemented using the draft-ietf-ipsec-ikev2-16.txt draft.

About IKE Policy Negotiation

To protect IKE negotiations, each IKE negotiation begins with a common (shared) IKE policy. An IKE policy defines a combination of security parameters to be used during the IKE negotiation. By default, no IKE policy is configured. You must create IKE policies at each peer. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how peers are authenticated. You can create multiple, prioritized policies at each peer to ensure that at least one policy will match a remote peer's policy.

You can configure the policy based on the encryption algorithm (DES, 3DES, or AES), the hash algorithm (SHA or MD5), and the DH group (1, 2, or 5). Each policy can contain a different combination of parameter values. A unique priority number identifies the configured policy. This number ranges from 1 (highest priority) to 255 (lowest priority). You can create multiple policies in a switch. If you need to connect to a remote peer, you must ascertain that at least one policy in the local switch contains the identical parameter values configured in the remote peer. If several policies have identical parameter configurations, the policy with the lowest number is selected.

Table 36-1 provides a list of allowed transform combinations.

Table 36-1 IKE Transform Configuration Parameters

Parameter	Accepted Values	Keyword	Default Value
encryption algorithm	56-bit DES-CBC	**des**	**3des**
	168-bit DES	**3des**	
	128-bit AES	**aes**	
hash algorithm	SHA-1 (HMAC variant)	**sha**	**sha**
	MD5 (HMAC variant)	**md5**	
authentication method	Preshared keys	Not configurable	Preshared keys
DH group identifier	768-bit DH	**1**	**1**
	1024-bit DH	**2**	
	1536-bit DH	**5**	

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

Note When you configure the hash algorithm, the corresponding HMAC version is used as the authentication algorithm.

When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the other peer's received policies. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is found when the two peers have the same encryption, hash algorithm, authentication algorithm, and DH group values. If a match is found, IKE completes the security negotiation and the IPsec SAs are created.

If an acceptable match is not found, IKE refuses negotiation and the IPsec data flows will not be established.

Configuring an IKE Policy

To configure the IKE negotiation parameters, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike domain ipsec` `switch(config-ike-ipsec)#`	Allows IPsec domains to be configured in this switch.
Step 3	`switch(config-ike-ipsec)# identity address`	Configures the identity mode for the IKE protocol to use the IP address (default).
	`switch(config-ike-ipsec)# identity hostname`	Configures the identity mode for the IKE protocol to use the fully-qualified domain name (FQDN). **Note** The FQDN is required for using RSA signatures for authentication.
	`switch(config-ike-ipsec)# no identity`	Revert to the default identity mode (**address**).

	Command	Purpose
Step 4	`switch(config-ike-ipsec)# key switch1 address 10.10.1.1`	Associates a preshared key with the IP address of a peer.
	`switch(config-ike-ipsec)# no key switch1 address 10.10.1.1`	Deletes the association of a preshared key and the IP address of a peer.
	`switch(config-ike-ipsec)# key switch1 hostname switch1.cisco.com`	Associates a preshared key with the FQDN of a peer. **Note** To use the FQDN, you must configure the switch name and domain name on the peer.
	`switch(config-ike-ipsec)# no key switch1 hostname switch1.cisco.com`	Deletes the association of a preshared key and the IP address of a peer.
Step 5	`switch(config-ike-ipsec)# policy 1` `switch(config-ike-ipsec-policy)#`	Specifies the policy to configure.
	`switch(config-ike-ipsec)# no policy 1`	Deletes the specified policy.
Step 6	`switch(config-ike-ipsec-policy)# encryption des`	Configures the encryption policy.
	`switch(config-ike-ipsec-policy)# no encryption des`	Defaults to 3DES encryption.
Step 7	`switch(config-ike-ipsec-policy)# group 5`	Configures the DH group.
	`switch(config-ike-ipsec-policy)# no group 5`	Defaults to DH group 1.
Step 8	`switch(config-ike-ipsec-policy)# hash md5`	Configures the hash algorithm.
	`switch(config-ike-ipsec-policy)# no hash md5`	Defaults to SHA.
Step 9	`switch(config-ike-ipsec-policy)# authentication pre-share`	Configures the authentication method to use the preshared key (default).
	`switch(config-ike-ipsec-policy)# authentication rsa-sig`	Configures the authentication method to use the RSA signature. **Note** To use RSA signatures for authentication you must configure identity authentication mode using the FQDN (see Step 3).
	`switch(config-ike-ipsec-policy)# no authentication`	Reverts to the default (**pre-share**).

Note When the authentication method is rsa-sig, make sure the identity hostname is configured for IKE because the IKE certificate has a subject name of the FQDN type.

Optional IKE Parameter Configuration

You can optionally configure the following parameters for the IKE feature:

- The lifetime association within each policy—The lifetime ranges from 600 to 86,400 seconds. The default is 86,400 seconds (equals one day). The lifetime association within each policy is configured when you are creating an IKE policy. See the "Configuring an IKE Policy" section on page 36-13.
- The keepalive time for each peer if you use IKEv2—The keepalive ranges from 120 to 86,400 seconds. The default is 3,600 seconds (equals one hour).
- The initiator version for each peer—IKE v1 or IKE v2 (default). Your choice of initiator version does not affect interoperability when the remote device initiates the negotiation. Configure this option if the peer device supports IKEv1 and you can play the initiator role for IKE with the specified device. Use the following considerations when configuring the initiator version with FCIP tunnels:
  - If the switches on both sides of an FCIP tunnel are running MDS SAN-OS Release 3.0(1) or later, you must configure initiator version IKEv1 on both sides of an FCIP tunnel to use only IKEv1. If one side of an FCIP tunnel is using IKEv1 and the other side is using IKEv2, the FCIP tunnel uses IKEv2.
  - If the switch on one side of an FCIP tunnel is running MDS SAN-OS Release 3.0(1) or later and the switch on the other side of the FCIP tunnel is running MDS SAN-OS Release 2.x, configuring IKEv1 on either side (or both) results in the FCIP tunnel using IKEv1.

Note Only IKE v1 is supported to build IPsec between 2.x and 3.x MDS switches.

Caution You may need to configure the initiator version even when the switch does not behave as an IKE initiator under normal circumstances. Always using this option guarantees a faster recovery of traffic flows in case of failures.

Tip The keepalive time only applies to IKEv2 peers and not to all peers.

Note When IPsec implementations in the host prefer to initiate the IPsec rekey, be sure to configure the IPsec lifetime value in the Cisco MDS switch to be higher than the lifetime value in the host.

This section includes the following topics:

Configuring the Lifetime Association for a Policy

To configure the lifetime association for each policy, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike domain ipsec` `switch(config-ike-ipsec)#`	Allows IPsec domains to be configured in this switch.
Step 3	`switch(config-ike-ipsec)# policy 1` `switch(config-ike-ipsec-policy)#`	Specified the policy to configure.
Step 4	`switch(config-ike-ipsec-policy) lifetime seconds 6000`	Configures a lifetime of 6,000 seconds.
	`switch(config-ike-ipsec-policy)# no lifetime seconds 6000`	Deletes the configured lifetime value and defaults to 86,400 seconds.

Configuring the Keepalive Time for a Peer

To configure the keepalive time for each peer, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike domain ipsec` `switch(config-ike-ipsec)#`	Allows IPsec domains to be configured in this switch.
Step 3	`switch(config-ike-ipsec)# keepalive 60000`	Configures the keepalive time for all peers to be 60,000 seconds.
	`switch(config-ike-ipsec)# no keepalive 60000`	Deletes the configured keepalive time and defaults to 3,600 seconds.

Configuring the Initiator Version

To configure the initiator version using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto ike domain ipsec` `switch(config-ike-ipsec)#`	Allows IPsec domains to be configured in this switch.
Step 3	`switch(config-ike-ipsec)# initiator version 1 address 10.10.10.1`	Configures the switch to use IKEv1 when initiating IKE with device 10.10.10.0 **Note** IKE supports IPv4 addresses, not IPv6 addresses.
	`switch(config-ike-ipsec)# no initiator version 1 address 10.10.10.1`	Defaults to IKEv2 for the specified device.
	`switch(config-ike-ipsec)# no initiator version 1`	Defaults to IKEv2 for all devices.

Clearing IKE Tunnels or Domains

If an IKE tunnel ID is not specified for the IKE configuration, you can clear all existing IKE domain connections by issuing the **clear crypto ike domain**. **ipsec sa** command in EXEC mode.

```
switch# clear crypto ike domain ipsec sa
```

Caution When you delete all the SAs within a specific IKEv2 tunnel, then that IKE tunnel is automatically deleted.

If an SA is specified for the IKE configuration, you can clear the specified IKE tunnel ID connection by issuing the **clear crypto ike domain ipsec sa** *IKE_tunnel-ID* command in EXEC mode.

```
switch# clear crypto ike domain ipsec sa 51
```

Caution When you delete the IKEv2 tunnel, the associated IPsec tunnel under that IKE tunnel is automatically deleted.

Refreshing SAs

Use the **crypto ike domain ipsec rekey IPv4-ACL**-*index* command to refresh the SAs after performing IKEv2 configuration changes.

Crypto IPv4-ACLs

IP access control lists (IPv4-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4 IP-ACLs restrict IP-related traffic based on the configured IP filters. See Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists" for details on creating and defining IPv4-ACLs.

In the context of crypto maps, IPv4-ACLs are different from regular IPv4-ACLs. Regular IPv4-ACLs determine what traffic to forward or block at an interface. For example, IPv4-ACLs can be created to protect all IP traffic between subnet A and subnet Y or Telnet traffic between host A and host B.

This section contains the following topics:

About Crypto IPv4-ACLs

Crypto IPv4-ACLs are used to define which IP traffic requires crypto protection and which traffic does not.

Crypto IPv4-ACLs associated with IPsec crypto map entries have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Indicate the data flow to be protected by the new SAs (specified by a single permit entry) when initiating negotiations for IPsec SAs.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether or not to accept requests for IPsec SAs on behalf of the requested data flows when processing IKE negotiation from the IPsec peer.

Tip If you want some traffic to receive one type of IPsec protection (for example, encryption only) and other traffic to receive a different type of IPsec protection (for example, both authentication and encryption), create two IPv4-ACLs. Use both IPv4-ACLs in different crypto maps to specify different IPsec policies.

Note IPsec does not support IPv6-ACLs.

Crypto IPv4-ACL Guidelines

Follow these guidelines when configuring IPv4-ACLs for the IPsec feature:

- The Cisco SAN-OS software only allows name-based IPv4-ACLs.
- When an IPv4-ACL is applied to a crypto map, the following options apply:
  - Permit—Applies the IPsec feature to the traffic.
  - Deny—Allows clear text (default).

  Note IKE traffic (UDP port 500) is implicitly transmitted in clear text.

- The IPsec feature only considers the source and destination IPv4 addresses and subnet masks, protocol, and single port number. There is no support for IPv6 in IPsec.

  Note The IPsec feature does not support port number ranges and ignores higher port number field, if specified.

- The permit option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.

- The deny option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto IPv4-ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different IPv4-ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPv4-ACL. Therefore, the IPv4-ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.
- Each IPv4-ACL filter assigned to the crypto map entry is equivalent to one security policy entry. The IPsec feature supports up to 120 security policy entries for each MPS-14/2 module and Cisco MDS 9216i Switch.
- In Figure 36-5, IPsec protection is applied to traffic between switch interface S0 (IPv4 address 10.0.0.1) and switch interface S1 (IPv4 address 20.0.0.2) as the data exits switch A's S0 interface enroute to switch interface S1. For traffic from 10.0.0.1 to 20.0.0.2, the IPv4-ACL entry on switch A is evaluated as follows:
  - source = IPv4 address 10.0.0.1
  - dest = IPv4 address 20.0.0.2

  For traffic from 20.0.0.2 to 10.0.0.1, that same IPv4-ACL entry on switch A is evaluated as follows:
  - source = IPv4 address 20.0.0.2
  - dest = IPv4 address 10.0.0.1

Figure 36-5 IPsec Processing of Crypto IPv4-ACLs

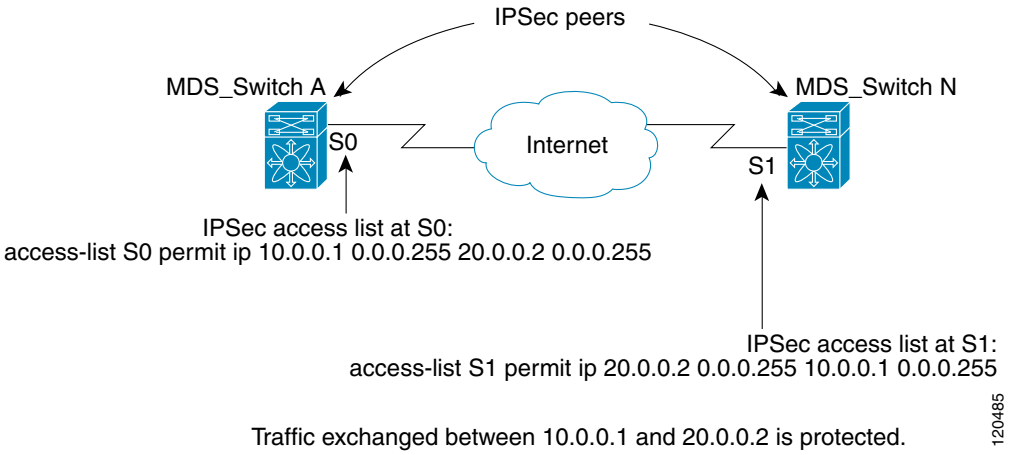

- If you configure multiple statements for a given crypto IPv4-ACL that is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto IPv4-ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched IPv4-ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto IPv4-ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.

- You can use the **show ip access-lists** command to view all IP-ACLs. The IP-ACLs used for traffic filtering purposes are also used for crypto.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures. The following example of a IPv4-ACL entry shows that the MDS switch IPv4 address is 10.10.10.50 and remote Microsoft host running encrypted iSCSI sessions is 10.10.10.16:

```
switch(config)# ip access-list aclmsiscsi2 permit tcp 10.10.10.50 0.0.0.0 range port
3260 3260 10.10.10.16 0.0.0.0
```

Mirror Image Crypto IPv4-ACLs

For every crypto IPv4-ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto IPv4-ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.

Tip The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.

Figure 36-6 shows some sample scenarios with and without mirror image IPv4-ACLs.

Figure 36-6 IPsec Processing of Mirror Image Configuration

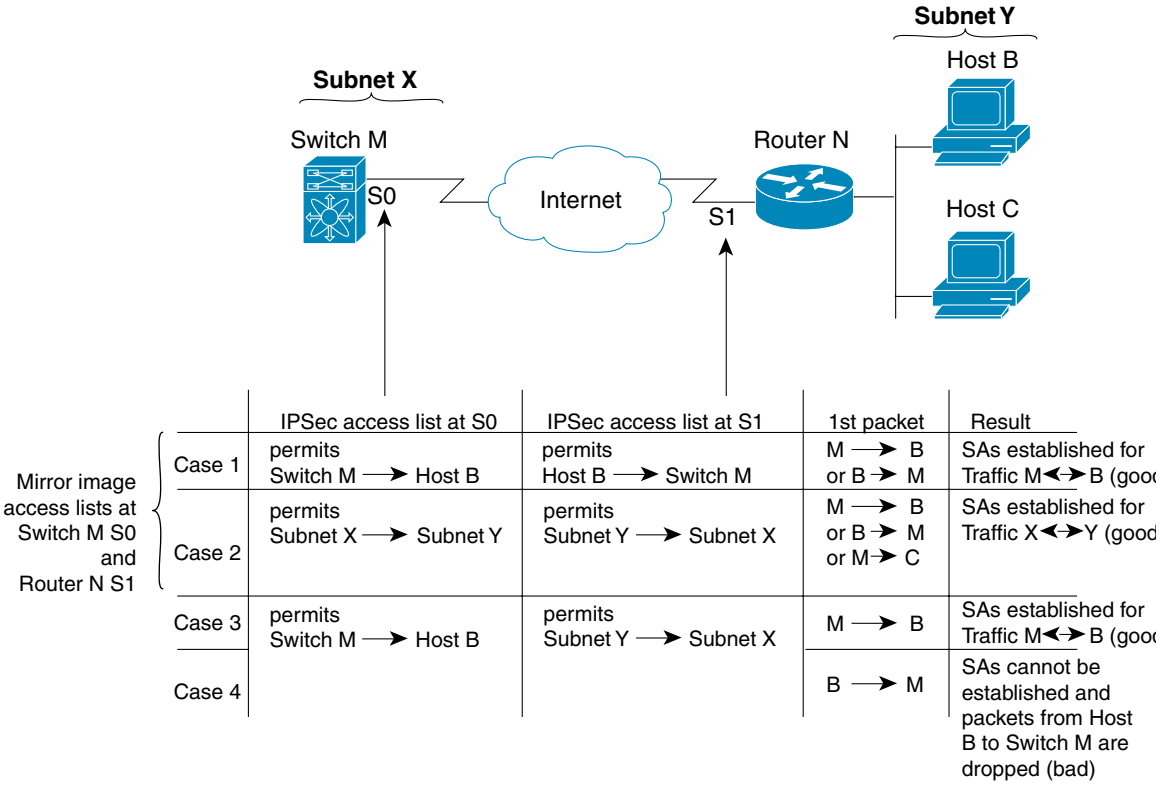

As Figure 36-6 indicates, IPsec SAs can be established as expected whenever the two peers' crypto IPv4-ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the IPv4-ACLs are not mirror images of each other. This can happen in the case when an entry in one peer's IPv4-ACL is a subset of an entry in the other peer's IPv4-ACL, such as shown in cases 3 and 4 of Figure 36-6. IPsec SA establishment is critical to IPsec. Without SAs, IPsec does not work, causing any packets matching the crypto IPv4-ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In case 4, an SA cannot be established because SAs are always requested according to the crypto IPv4-ACLs at the initiating packet's end. In case 4, router N requests that all traffic between subnet X and subnet Y be protected, but this is a superset of the specific flows permitted by the crypto IPv4-ACL at switch M so the request is not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto IPv4-ACL at router N.

Because of the complexities introduced when crypto IPv4-ACLs are not configured as mirror images at peer IPsec devices, we strongly encourage you to use mirror image crypto IPv4-ACLs.

The any Keyword in Crypto IPv4-ACLs

Tip We recommend that you configure mirror image crypto IPv4-ACLs for use by IPsec and that you avoid using the **any** option.

The **any** keyword in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface. This configuration can cause multicast traffic to fail.

The **permit any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use **any** in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

Creating Crypto IPv4-ACLs

To create IPv4-ACLs, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ip access-list List1 permit ip 10.1.1.100 0.0.0.255 11.1.1.100 0.0.0.255`	Permits all IP traffic from and to the specified networks.

Note The **show ip access-list** command does not display the crypto map entries. Use the **show crypto map** command to display the associated entries.

Add permit and deny statements as appropriate (see Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists,"). Each permit and deny specifies conditions to determine which IP packets must be protected.

About Transform Sets in IPsec

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations.

Tip If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.

Note When you enable IPsec, the Cisco SAN-OS software automatically creates a default transform set (ipsec_default_tranform_set) using AES-128 encryption and SHA-1 authentication algorithms.

Table 36-2 provides a list of allowed transform combinations for IPsec.

Table 36-2 IPsec Transform Configuration Parameters

Parameter	Accepted Values	Keyword
encryption algorithm	56-bit DES-CBC	**esp-des**
	168-bit DES	**esp-3des**
	128-bit AES-CBC	**esp-aes 128**
	128-bit AES-CTR[1]	**esp-aes 128 ctr**
	256-bit AES-CBC	**esp-aes 256**
	256-bit AES-CTR[1]	**esp-aes 256 ctr**
hash/authentication algorithm[1] (optional)	SHA-1 (HMAC variant)	**esp-sha1-hmac**
	MD5 (HMAC variant)	**esp-md5-hmac**
	AES-XCBC-MAC	**esp-aes-xcbc-mac**

1. If you configure the AES counter (CTR) mode, you must also configure the authentication algorithm.

Note The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

Configuring Transform Sets

To configure transform sets, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto transform-set domain ipsec test esp-3des esp-md5-hmac`	Configures a transform set called test specifying the 3DES encryption algorithm and the MD5 authentication algorithm. Refer to Table 36-2 to verify the allowed transform combinations.
	`switch(config)# no crypto transform-set domain ipsec test esp-3des esp-md5-hmac`	Deletes the applied transform set.
	`switch(config)# crypto transform-set domain ipsec test esp-3des`	Configures a transform set called test specifying the 3DES encryption algorithm. In this case, the default no authentication is performed.
	`switch(config)# no crypto transform-set domain ipsec test esp-3des`	Deletes the applied transform set.

About Crypto Map Entries

Once you have created the crypto IPv4-ACLs and transform sets, you can create crypto map entries that combine the various parts of the IPsec SA, including the following:

- The traffic to be protected by IPsec (per the crypto IPv4-ACL). A crypto map set can contain multiple entries, each with a different IPv4-ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets).
- Other parameters to define an IPsec SA.

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A security policy database (SPD) is created for that interface.
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decides whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto IPv4-ACLs (for example, mirror image IPv4-ACLs). If the responding peer entry is in the local crypto, the IPv4-ACL must be permitted by the peer's crypto IPv4-ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.
- If you create more than one crypto map entry for a given interface, use the `seq-num` of each map entry to rank the map entries: the lower the `seq-num`, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common, where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular IPv4-ACL, the corresponding crypto map entry is tagged, and the connections are established.

Crypto Map Configuration Guidelines

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.
- Only one IPv4-ACL is allowed for each crypto map entry (the IPv4-ACL itself can have multiple permit or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the auto-peer option to dynamically configure the peer.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.

Creating Crypto Map Entries

Note If the peer IP address specified in the crypto map entry is a VRRP IP address on a remote Cisco MDS switch, ensure that the IP address is created using the **secondary** option (see the "Adding Virtual Router IP Addresses" section on page 43-20).

To create mandatory crypto map entries, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto map domain ipsec SampleMap 31` `ips-hac1(config-crypto-map-ip)#`	Place you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number.
	`switch(config)# no crypto map domain ipsec SampleMap 3`	Deletes the specified crypto map entry.
	`switch(config)# no crypto map domain ipsec SampleMap`	Deletes the entire crypto map set called SampleMap.
Step 3	`switch(config-crypto-map-ip)# match address SampleAcl`	Names an ACL to determine which traffic should be protected and not protected by IPsec in the context of this crypto map entry.
	`switch(config-crypto-map-ip)# no match address SampleAcl`	Deletes the matched address.
Step 4	`switch(config-crypto-map-ip)# set peer 10.1.1.1`	Configures a specific peer IPv4 address. **Note** IKE only supports IPv4 addresses, not IPv6 addresses.
Step 5	`switch(config-crypto-map-ip)# no set peer 10.1.1.1`	Deletes the configured peer.
Step 6	`switch(config-crypto-map-ip)# set transform-set SampleTransform1 SampleTransmfor2`	Specifies which transform sets are allowed for the specified crypto map entry or entries. List multiple transform sets in order of priority (highest priority first).
	`switch(config-(crypto-map-ip))# no set transform-set`	Deletes the association of all transform sets (regardless of you specifying a transform set name).

About SA Lifetime Negotiation

You can override the global lifetime values (size and time) by configuring an SA-specific lifetime value.

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.

See the "Global Lifetime Values" section on page 36-29 for more information on global lifetime values.

Setting the SA Lifetime

To set the SA lifetime for a specified crypto map entry, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto map domain ipsec SampleMap 31` `switch(config-crypto-map-ip)#`	Enters crypto map configuration submode for the entry named SampleMap with 31 as its sequence number.
Step 3	`switch(config-crypto-map-ip)# set security-association lifetime seconds 8640`	Specifies an SA lifetime for this crypto map entry using different IPsec SA lifetimes than the global lifetimes for the crypto map entry.
	`switch(config-crypto-map-ip)# no set security-association lifetime seconds 8640`	Deletes the entry-specific configuration and reverts to the global settings.
Step 4	`switch(config-crypto-map-ip)# set security-association lifetime kilobytes 2560`	Configures the traffic-volume lifetime for this SA in kilobytes. The lifetime ranges from 2560 to 2147483647 kilobytes.
	`switch(config-crypto-map-ip)# set security-association lifetime gigabytes 4000`	Configures the traffic-volume lifetime for this SA to time out after the specified amount of traffic (in gigabytes) have passed through the FCIP link using the SA. The lifetime ranges from 1 to 4095 gigabytes.
	`switch(config-crypto-map-ip)# set security-association lifetime megabytes 5000`	Configures the traffic-volume lifetime for this SA in megabytes. The lifetime ranges from 3 to 4193280 megabytes.
	`switch(config-crypto-map-ip)# no set security-association lifetime megabytes`	Reverts to the global settings.

About the AutoPeer Option

Setting the peer address as **auto-peer** in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be set up at each of the endpoints in the subnet specified by the crypto map's IPv4-ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Figure 36-7 shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to set up SAs with the switch. Each host will set up its own SA, but will share the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.

See the "Sample iSCSI Configuration" section on page 36-40 for more details.

Figure 36-7 iSCSI with End-to-End IPsec Using the auto-peer Option

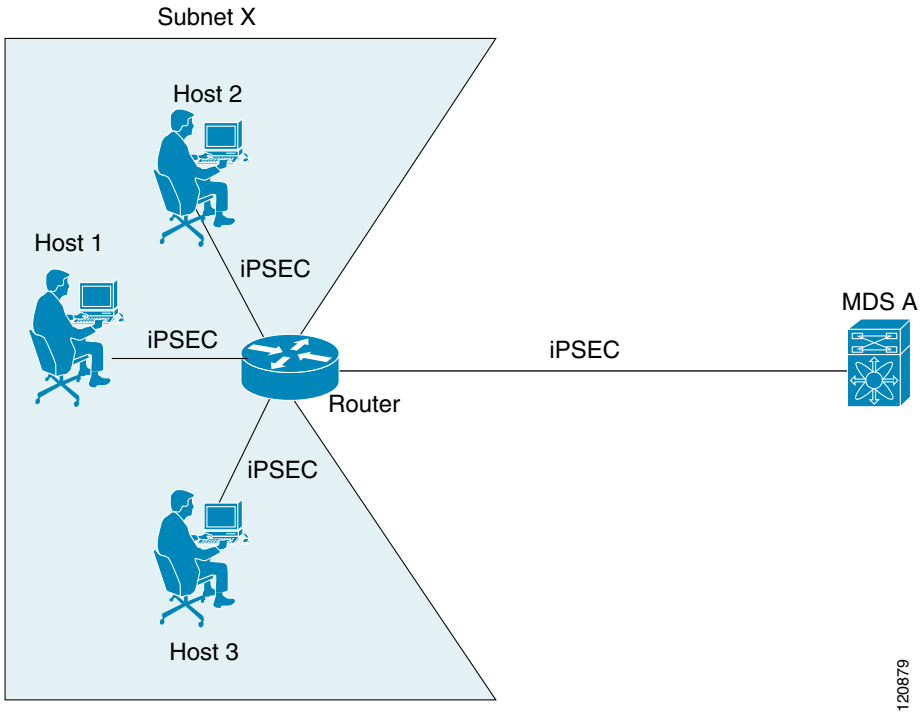

Configuring the AutoPeer Option

To configure the auto-peer option, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto map domain ipsec SampleMap 31` `ips-hac1(config-crypto-map-ip)#`	Places you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number.
Step 3	`switch(config-crypto-map-ip)# set peer auto-peer`	Directs the software to select (during the SA setup) the destination peer IP address dynamically.
	`switch(config-crypto-map-ip)# no set peer auto-peer`	Deletes the auto-peer configuration.

About Perfect Forward Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forward secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of the DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

Configuring Perfect Forward Secrecy

To configure the PFS value, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto map domain ipsec SampleMap 31` `ips-hac1(config-crypto-map-ip)#`	Places you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number.
Step 3	`switch(config-crypto-map-ip)# set pfs group 2`	Specifies that IPsec should ask for PFS when requesting new SAs for this crypto map entry, or should demand PFS in requests received from the IPsec peer.
	`switch(config-crypto-map-ip)# no set pfs`	Deletes the configured DH group and reverts to the factory default of disabling PFS.

About Crypto Map Set Interface Application

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of the traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

Applying a Crypto Map Set

To apply a crypto map set to an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 4/1` `switch(config-if)#`	Selects the required Gigabit Ethernet interface (and subinterface, if required) to which the IPsec crypto map is to be applied.

	Command	Purpose
Step 3	`switch(config-if)# crypto map domain ipsec cm10`	Applies the crypto map set to the selected interface.
Step 4	`switch(config-if)# no crypto map domain ipsec`	Deletes the crypto map that is currently applied to this interface.

IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be reestablished with the changed configuration. If the switch is actively processing IPsec traffic, it is desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.

Caution Using the **clear crypto sa** command without parameters will clear out the full SA database, which will clear out active security sessions. You may also specify the peer, map, or entry keywords to clear out only a subset of the SA database.

Tip You can obtain the SA index from the output of the **show crypto sa domain interface gigabitethernet** *slot/port* command.

Use the **clear crypto sa** command to clear all or part of the SA database.

```
switch# clear crypto sa domain ipsec interface gigabitethernet 2/1 inbound sa 1
```

Global Lifetime Values

If you have not configured a lifetime in the crypto map entry, the global lifetime values are used when negotiating new IPsec SAs.

You can configure two lifetimes: timed or traffic-volume. An SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 450 GB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKEv1 to set up IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKEv2 to set up IPsec SAs, the SAs on each end have their own lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold of the existing SA is reached to ensure that negotiation completes before the existing SA expires.

The new SA is negotiated when one of the following thresholds is reached (whichever comes first):

- 30 seconds before the lifetime expires or
- Approximately 10% of the lifetime in bytes remain

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

To configure global SA lifetimes, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# crypto global domain ipsec security-association lifetime seconds 86400`	Configures the global timed lifetime for IPsec SAs to time out after the specified number of seconds have passed. The global lifetime ranges from 120 to 86400 seconds.
	`switch(config)# no crypto global domain ipsec security-association lifetime seconds 86400`	Reverts to the factory default of 3,600 seconds.
Step 3	`switch(config)# crypto global domain ipsec security-association lifetime gigabytes 4000`	Configures the global traffic-volume lifetime for IPsec SAs to time out after the specified amount of traffic (in gigabytes) has passed through the FCIP link using the SA. The global lifetime ranges from 1 to 4095 gigabytes.
	`switch(config)# crypto global domain ipsec security-association lifetime kilobytes 2560`	Configures the global traffic-volume lifetime in kilobytes. The global lifetime ranges from 2560 to 2147483647 kilobytes.
	`switch(config)# crypto global domain ipsec security-association lifetime megabytes 5000`	Configures the global traffic-volume lifetime in megabytes. The global lifetime ranges from 3 to 4193280 megabytes.
	`switch(config)# no crypto global domain ipsec security-association lifetime megabytes`	Reverts to the factory default of 450 GB regardless of what value is currently configured.

Displaying IKE Configurations

You can verify the IKE information by using the **show** set of commands. See Examples 36-1 to 36-5.

Example 36-1 Displays the Parameters Configured for Each IKE Policy

```
switch# show crypto ike domain ipsec
keepalive 60000
```

Example 36-2 Displays the Initiator Configuration

```
switch# show crypto ike domain ipsec initiator
initiator version 1 address 1.1.1.1
initiator version 1 address 1.1.1.2
```

Example 36-3 Displays the Key Configuration

```
switch# show crypto ike domain ipsec key
key abcdefgh address 1.1.1.1
key bcdefghi address 1.1.2.1
```

Example 36-4 Displays the Currently Established Policies for IKE

```
switch# show crypto ike domain ipsec policy 1
Priority 1, auth pre-shared, lifetime 6000 secs, encryption 3des, hash md5, DH group 5
Priority 3, auth pre-shared, lifetime 86300 secs, encryption aes, hash sha1, DH group 1
```

Example 36-5 Displays the Currently Established SAs for IKE

```
switch# show crypto ike domain ipsec sa
Tunn    Local Addr          Remote Addr        Encr     Hash    Auth Method     Lifetime
------------------------------------------------------------------------------------------
1*      172.22.31.165[500]  172.22.31.166[500] 3des     sha1    preshared key   86400
2       172.22.91.174[500]  172.22.91.173[500] 3des     sha1    preshared key   86400
------------------------------------------------------------------------------------------
NOTE: tunnel id ended with * indicates an IKEv1 tunnel
```

Displaying IPsec Configurations

You can verify the IPsec information by using the **show** set of commands. See Examples 36-6 to 36-19.

Example 36-6 Displays Information for the Specified ACL

```
switch# show ip access-list acl10
ip access-list acl10 permit ip 10.10.10.0 0.0.0.255 10.10.10.0 0.0.0.255 (0 matches)
```

In Example 36-6, the display output match is only displayed of an interface (not the crypto map) meets this criteria.

Example 36-7 Displays the Transform Set Configuration

```
switch# show crypto transform-set domain ipsec
Transform set: 3des-md5 {esp-3des esp-md5-hmac}
    will negotiate {tunnel}
```

```
Transform set: des-md5 {esp-des esp-md5-hmac}
    will negotiate {tunnel}
Transform set: test {esp-aes-128-cbc esp-md5-hmac}
    will negotiate {tunnel}
```

Example 36-8 Displays All Configured Crypto Maps

```
switch# show crypto map domain ipsec
Crypto Map "cm10" 1 ipsec
        Peer = Auto Peer
        IP ACL = acl10
            permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
        Transform-sets: 3des-md5, des-md5,
        Security Association Lifetime: 4500 megabytes/3600 seconds
        PFS (Y/N): N
        Interface using crypto map set cm10:
            GigabitEthernet4/1
Crypto Map "cm100" 1 ipsec
        Peer = Auto Peer
        IP ACL = acl100
            permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
        Transform-sets: 3des-md5, des-md5,
        Security Association Lifetime: 4500 megabytes/3600 seconds
        PFS (Y/N): N
        Interface using crypto map set cm100:
            GigabitEthernet4/2
```

Example 36-9 Displays the Crypto Map Information for a Specific Interface

```
switch# show crypto map domain ipsec interface gigabitethernet 4/1
Crypto Map "cm10" 1 ipsec
        Peer = Auto Peer
        IP ACL = acl10
            permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
        Transform-sets: 3des-md5, des-md5,
        Security Association Lifetime: 4500 megabytes/120 seconds
        PFS (Y/N): N
        Interface using crypto map set cm10:
            GigabitEthernet4/1
```

Example 36-10 Displays the Specified Crypto Map Information

```
switch# show crypto map domain ipsec tag cm100
Crypto Map "cm100" 1 ipsec
        Peer = Auto Peer
        IP ACL = acl100
            permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
        Transform-sets: 3des-md5, des-md5,
        Security Association Lifetime: 4500 megabytes/120 seconds
        PFS (Y/N): N
        Interface using crypto map set cm100:
            GigabitEthernet4/2
```

Example 36-11 Displays SA Association for the Specified Interface

```
switch# show crypto sad domain ipsec interface gigabitethernet 4/1
interface: GigabitEthernet4/1
```

```
    Crypto map tag: cm10, local addr. 10.10.10.1
    protected network:
    local ident (addr/mask): (10.10.10.0/255.255.255.0)
    remote  ident (addr/mask): (10.10.10.4/255.255.255.255)
    current_peer: 10.10.10.4
      local crypto endpt.: 10.10.10.1, remote crypto endpt.: 10.10.10.4
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
     current outbound spi: 0x30e000f (51249167), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
     current inbound spi: 0x30e0000 (51249152), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
```

Example 36-12 Displays All SA Associations

```
switch# show crypto sad domain ipsec
interface: GigabitEthernet4/1
    Crypto map tag: cm10, local addr. 10.10.10.1
    protected network:
    local  ident (addr/mask): (10.10.10.0/255.255.255.0)
    remote  ident (addr/mask): (10.10.10.4/255.255.255.255)
    current_peer: 10.10.10.4
      local crypto endpt.: 10.10.10.1, remote crypto endpt.: 10.10.10.4
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
     current outbound spi: 0x30e000f (51249167), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
     current inbound spi: 0x30e0000 (51249152), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
```

Example 36-13 Displays Information About the Policy Database

```
switch# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet4/1, direction: Both
#   0:      deny  udp any port eq 500 any
#   1:      deny  udp any any port eq 500
#   2:      permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
#  63:      deny  ip any any
Policy Database for interface: GigabitEthernet4/2, direction: Both
#   0:      deny  udp any port eq 500 any <----------------------UDP default entry
#   1:      deny  udp any any port eq 500 <---------------------- UDP default entry
#   3:      permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
#  63:      deny  ip any any <------------------------------------- Clear text default entry
```

Example 36-14 Displays SPD Information for a Specific Interface

```
switch# show crypto spd domain ipsec interface gigabitethernet 4/2
Policy Database for interface: GigabitEthernet3/1, direction: Both
#   0:      deny  udp any port eq 500 any
#   1:      deny  udp any any port eq 500
#   2:      permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
# 127:      deny  ip any any
```

Example 36-15 Displays Detailed iSCSI Session Information for a Specific Interface

```
switch# show iscsi session detail
Initiator iqn.1987-05.com.cisco:01.9f39f09c7468 (ips-host16.cisco.com)
    Initiator ip addr (s): 10.10.10.5
  Session #1 (index 24)
    Discovery session, ISID 00023d000001, Status active

  Session #2 (index 25)
    Target ibm1
    VSAN 1, ISID 00023d000001, TSIH 0, Status active, no reservation
    Type Normal, ExpCmdSN 42, MaxCmdSN 57, Barrier 0
    MaxBurstSize 0, MaxConn 1, DataPDUInOrder Yes
    DataSeqInOrder Yes, InitialR2T Yes, ImmediateData No
    Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 41, Response: 41
      Bytes: TX: 21388, RX: 0
    Number of connection: 1
    Connection #1
      iSCSI session is protected by IPSec <----------The iSCSI session protection status
      Local IP address: 10.10.10.4, Peer IP address: 10.10.10.5
      CID 0, State: Full-Feature
      StatSN 43, ExpStatSN 0
      MaxRecvDSLength 131072, our_MaxRecvDSLength 262144
      CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
      AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
      Version Min: 0, Max: 0
      FC target: Up, Reorder PDU: No, Marker send: No (int 0)
      Received MaxRecvDSLen key: Yes
```

Example 36-16 Displays FCIP Information for a Specific Interface

```
switch# show interface fcip 1
fcip1 is trunking
    Hardware is GigabitEthernet
    Port WWN is 20:50:00:0d:ec:08:6c:c0
    Peer port WWN is 20:10:00:05:30:00:a7:9e
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 1 Gbps
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    Using Profile id 1  (interface GigabitEthernet2/1)
    Peer Information
      Peer Internet address is 10.10.11.1 and port is 3225
    FCIP tunnel is protected by IPSec <----------The FCIP tunnel protection status
    Write acceleration mode is off
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is 256 KBytes
    IP Compression is disabled
    Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
    TCP Connection Information
      2 Active TCP connections
```

```
        Control connection: Local 10.10.11.2:3225, Remote 10.10.11.1:65520
        Data connection: Local 10.10.11.2:3225, Remote 10.10.11.1:65522
      2 Attempts for active connections, 0 close of connections
    TCP Parameters
      Path MTU 1400 bytes
      Current retransmission timeout is 200 ms
      Round trip time: Smoothed 2 ms, Variance: 1
      Advertized window: Current: 124 KB, Maximum: 124 KB, Scale: 6
      Peer receive window: Current: 123 KB, Maximum: 123 KB, Scale: 6
      Congestion window: Current: 53 KB, Slow start threshold: 48 KB
      Current Send Buffer Size: 124 KB, Requested Send Buffer Size: 0 KB
      CWM Burst Size: 50 KB
    5 minutes input rate 128138888 bits/sec, 16017361 bytes/sec, 7937 frames/sec
    5 minutes output rate 179275536 bits/sec, 22409442 bytes/sec, 46481 frames/sec
      10457037 frames input, 21095415496 bytes
         308 Class F frames input, 32920 bytes
         10456729 Class 2/3 frames input, 21095382576 bytes
         9907495 Reass frames
         0 Error frames timestamp error 0
      63792101 frames output, 30250403864 bytes
         472 Class F frames output, 46816 bytes
         63791629 Class 2/3 frames output, 30250357048 bytes
         0 Error frames
```

Example 36-17 Displays the Global IPsec Statistics for the Switch

```
switch# show crypto global domain ipsec
IPSec global statistics:
        Number of crypto map sets: 3
        IKE transaction stats: 0 num, 256 max
        Inbound SA stats: 0 num
        Outbound SA stats: 0 num
```

Example 36-18 Displays the IPsec Statistics for the Specified Interface

```
switch# show crypto global domain ipsec interface gigabitethernet 3/1
IPSec interface statistics:
        IKE transaction stats: 0 num
        Inbound SA stats: 0 num, 512 max
        Outbound SA stats: 0 num, 512 max
```

Example 36-19 Displays the Global SA Lifetime Values

```
switch# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 450 gigabytes/3600 seconds
```

Sample FCIP Configuration

Figure 36-8 focuses on implementing IPsec for one FCIP link (Tunnel 2). Tunnel 2 carries encrypted data between MDS A and MDS C.

Figure 36-8 IP Security Usage in an FCIP Scenario

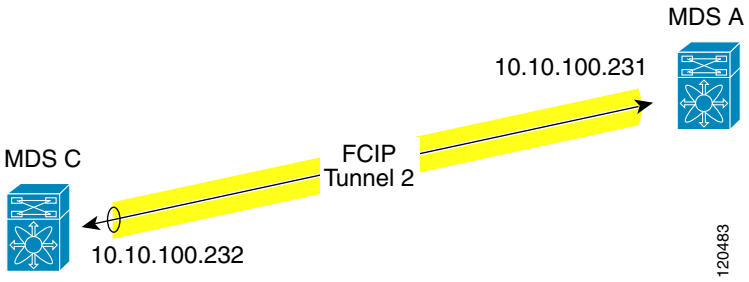

To configure IPsec for the FCIP scenario shown in Figure 36-8, follow these steps:

Step 1 Enable IKE and IPsec in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# crypto ike enable
sw10.1.1.100(config)# crypto ipsec enable
```

Step 2 Configure IKE in Switch MDS A.

```
sw10.1.1.100(config)# crypto ike domain ipsec
sw10.1.1.100(config-ike-ipsec)# key ctct address 10.10.100.232
sw10.1.1.100(config-ike-ipsec)# policy 1
sw10.1.1.100(config-ike-ipsec-policy)# encryption 3des
sw10.1.1.100(config-ike-ipsec-policy)# hash md5
sw10.1.1.100(config-ike-ipsec-policy)# end
sw10.1.1.100#
```

Step 3 Configure the ACLs in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# ip access-list acl1 permit tcp 10.10.100.231 0.0.0.0 range port 3260
3260 10.10.100.232 0.0.0.0
```

Step 4 Configure the transform set in Switch MDS A.

```
sw10.1.1.100(config)# crypto transform-set domain ipsec tfs-02 esp-aes 128 esp-sha1-hmac
```

Step 5 Configure the crypto map in Switch MDS A.

```
sw10.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw10.1.1.100(config-crypto-map-ip)# match address acl1
sw10.1.1.100(config-crypto-map-ip)# set peer 10.10.100.232
sw10.1.1.100(config-crypto-map-ip)# set transform-set tfs-02
sw10.1.1.100(config-crypto-map-ip)# set security-association lifetime seconds 120
sw10.1.1.100(config-crypto-map-ip)# set security-association lifetime gigabytes 3000
sw10.1.1.100(config-crypto-map-ip)# set pfs group5
sw10.1.1.100(config-crypto-map-ip)# end
sw10.1.1.100#
```

Step 6 Bind the interface to the crypto map set in Switch MDS A.

```
sw10.1.1.100# conf t
```

```
sw10.1.1.100(config)# int gigabitethernet 7/1
sw10.1.1.100(config-if)# ip addr 10.10.100.231 255.255.255.0
sw10.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# exit
sw10.1.1.100(config)#
```

Step 7 Configure FCIP in Switch MDS A.

```
sw10.1.1.100(config)# fcip enable
sw10.1.1.100(config)# fcip profile 2
sw10.1.1.100(config-profile)# ip address 10.10.100.231
sw10.1.1.100(config-profile)# int fcip 2
sw10.1.1.100(config-if)# peer-info ipaddr 10.10.100.232
sw10.1.1.100(config-if)# use-profile 2
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# end
sw10.1.1.100#
```

Step 8 Verify the configuration in Switch MDS A.

```
sw10.1.1.100# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 4500 megabytes/3600 seconds

sw10.1.1.100# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
        Peer = 10.10.100.232
        IP ACL = acl1
            permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
        Transform-sets: tfs-02,
        Security Association Lifetime: 3000 gigabytes/120 seconds
        PFS (Y/N): Y
         PFS Group: group5
Interface using crypto map set cmap-01:
    GigabitEthernet7/1

sw10.1.1.100# show crypto transform-set domain ipsec
Transform set: tfs-02 {esp-aes 128 esp-sha1-hmac}
    will negotiate {tunnel}

sw10.1.1.100# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet7/1, direction: Both
#   0:      deny  udp any port eq 500 any
#   1:      deny  udp any any port eq 500
#   2:      permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
#  63:      deny  ip any any

sw10.1.1.100# show crypto ike domain ipsec
keepalive 3600

sw10.1.1.100# show crypto ike domain ipsec key
key ctct address 10.10.100.232

sw10.1.1.100# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH group 1
```

Step 9 Enable IKE and IPsec in Switch MDS C.

```
sw11.1.1.100# conf t
sw11.1.1.100(config)# crypto ike enable
sw11.1.1.100(config)# crypto ipsec enable
```

Step 10 Configure IKE in Switch MDS C.

```
sw11.1.1.100(config)# crypto ike domain ipsec
```

```
sw11.1.1.100(config-ike-ipsec)# key ctct address 10.10.100.231
sw11.1.1.100(config-ike-ipsec)# policy 1
sw11.1.1.100(config-ike-ipsec-policy)# encryption 3des
sw11.1.1.100(config-ike-ipsec-policy)# hash md5
sw11.1.1.100(config-ike-ipsec-policy)# exit
sw11.1.1.100(config-ike-ipsec)# end
sw11.1.1.100#
```

Step 11 Configure the ACLs in Switch MDS C.

```
sw11.1.1.100# conf t
sw11.1.1.100(config)# ip access-list acl1 permit ip 10.10.100.232 0.0.0.0 10.10.100.231
0.0.0.0
```

Step 12 Configure the transform set in Switch MDS C.

```
sw11.1.1.100(config)# crypto transform-set domain ipsec tfs-02 esp-aes 128 esp-sha1-hmac
```

Step 13 Configure the crypto map in Switch MDS C.

```
sw11.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw11.1.1.100(config-crypto-map-ip)# match address acl1
sw11.1.1.100(config-crypto-map-ip)# set peer 10.10.100.231
sw11.1.1.100(config-crypto-map-ip)# set transform-set tfs-02
sw11.1.1.100(config-crypto-map-ip)# set security-association lifetime seconds 120
sw11.1.1.100(config-crypto-map-ip)# set security-association lifetime gigabytes 3000
sw11.1.1.100(config-crypto-map-ip)# set pfs group5
sw11.1.1.100(config-crypto-map-ip)# exit
sw11.1.1.100(config)#
```

Step 14 Bind the interface to the crypto map set in Switch MDS C.

```
sw11.1.1.100(config)# int gigabitethernet 1/2
sw11.1.1.100(config-if)# ip addr 10.10.100.232 255.255.255.0
sw11.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw11.1.1.100(config-if)# no shut
sw11.1.1.100(config-if)# exit
sw11.1.1.100(config)#
```

Step 15 Configure FCIP in Switch MDS C.

```
sw11.1.1.100(config)# fcip enable
sw11.1.1.100(config)# fcip profile 2
sw11.1.1.100(config-profile)# ip address 10.10.100.232
sw11.1.1.100(config-profile)# int fcip 2
sw11.1.1.100(config-if)# peer-info ipaddr 10.10.100.231
sw11.1.1.100(config-if)# use-profile 2
sw11.1.1.100(config-if)# no shut
sw11.1.1.100(config-if)# exit
sw11.1.1.100(config)# exit
```

Step 16 Verify the configuration in Switch MDS C.

```
sw11.1.1.100# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 4500 megabytes/3600 seconds

sw11.1.1.100# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
        Peer = 10.10.100.231
        IP ACL = acl1
            permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
        Transform-sets: tfs-02,
        Security Association Lifetime: 3000 gigabytes/120 seconds
        PFS (Y/N): Y
         PFS Group: group5
```

```
Interface using crypto map set cmap-01:
    GigabitEthernet1/2

sw11.1.1.100# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet1/2, direction: Both
#   0:      deny  udp any port eq 500 any
#   1:      deny  udp any any port eq 500
#   2:      permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
#  63:      deny  ip any any

sw11.1.1.100# show crypto sad domain ipsec
interface: GigabitEthernet1/2
    Crypto map tag: cmap-01, local addr. 10.10.100.232
    protected network:
    local  ident (addr/mask): (10.10.100.232/255.255.255.255)
    remote  ident (addr/mask): (10.10.100.231/255.255.255.255)
    current_peer: 10.10.100.231
      local crypto endpt.: 10.10.100.232, remote crypto endpt.: 10.10.100.231
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
     current outbound spi: 0x38f96001 (955867137), index: 29
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000
     current inbound spi: 0x900b011 (151040017), index: 16
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000

sw11.1.1.100# show crypto transform-set domain ipsec
Transform set: tfs-02 {esp-aes 128 esp-sha1-hmac}
    will negotiate {tunnel}

sw11.1.1.100# show crypto ike domain ipsec
keepalive 3600

sw11.1.1.100# show crypto ike domain ipsec key

key ctct address 10.10.100.231

sw11.1.1.100# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH
group 1

sw11.1.1.100# show crypto ike domain ipsec sa
Tunn    Local Addr          Remote Addr         Encr    Hash    Auth Method     Lifetime
------------------------------------------------------------------------------------------
1*      10.10.100.232[500]  10.10.100.231[500]  3des    md5     preshared key   86300
------------------------------------------------------------------------------------------
NOTE: tunnel id ended with * indicates an IKEv1 tunnel
```

Step 17 Verify the configuration in Switch MDS A.

```
sw10.1.1.100# show crypto sad domain ipsec
interface: GigabitEthernet7/1
    Crypto map tag: cmap-01, local addr. 10.10.100.231
    protected network:
    local  ident (addr/mask): (10.10.100.231/255.255.255.255)
    remote  ident (addr/mask): (10.10.100.232/255.255.255.255)
    current_peer: 10.10.100.232
      local crypto endpt.: 10.10.100.231, remote crypto endpt.: 10.10.100.232
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
     current outbound spi: 0x900b01e (151040030), index: 10
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000
     current inbound spi: 0x38fe700e (956198926), index: 13
      lifetimes in seconds:: 120
```

```
      lifetimes in bytes:: 3221225472000

sw10.1.1.100# show crypto ike domain ipsec sa
Tunn Local Addr          Remote Addr         Encr  Hash  Auth Method   Lifetime
------------------------------------------------------------------------------
   1 10.10.100.231[500]  10.10.100.232[500]  3des  md5   preshared key    86300
```

You have now configured IPsec in both switches MDS A and MDS C.

Sample iSCSI Configuration

Figure 36-9 focuses on the iSCSI session between MDS A and the hosts in subnet 12.12.1/24. Using the **auto-peer** option, when any host from the subnet 12.12.1.0/24 tries to connect to the MDS switch's Gigabit Ethernet port 7/1, an SA is created between the hosts and the MDS switch. With auto-peer, only one crypto map is necessary to create SAs for all the hosts in the same subnet. Without auto-peer, you need one crypto map entry per host.

Figure 36-9 iSCSI with End-to-End IPsec

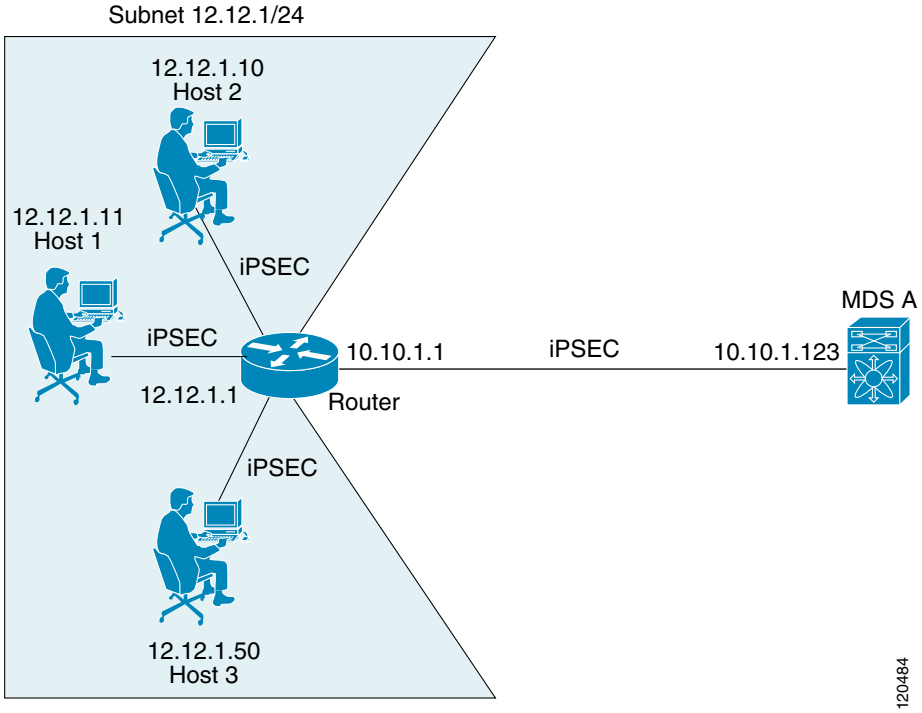

To configure IPsec for the iSCSI scenario shown in Figure 36-9, follow these steps:

Step 1 Configure the ACLs in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# ip access-list acl1 permit tcp 10.10.1.0 0.0.0.255 range port 3260
3260 12.12.1.0 0.0.0.255
```

Step 2 Configure the transform set in Switch MDS A.

```
sw10.1.1.100(config)# crypto transform-set domain ipsec tfs-01 esp-3des esp-md5-hmac
```

Step 3 Configure the crypto map in Switch MDS A.

```
sw10.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw10.1.1.100(config-crypto-map-ip)# match address acl1
sw10.1.1.100(config-crypto-map-ip)# set peer auto-peer
sw10.1.1.100(config-crypto-map-ip)# set transform-set tfs-01
sw10.1.1.100(config-crypto-map-ip)# end
sw10.1.1.100#
```

Step 4 Bind the interface to the crypto map set in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# int gigabitethernet 7/1
sw10.1.1.100(config-if)# ip address 10.10.1.123 255.255.255.0
sw10.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# end
sw10.1.1.100#
```

You have now configured IPsec in MDS A using the Cisco MDS IPsec and iSCSI features.

Default Settings

Table 36-3 lists the default settings for IKE parameters.

Table 36-3 Default IKE Parameters

Parameters	Default
IKE	Disabled.
IKE version	IKE version 2.
IKE encryption algorithm	3DES.
IKE hash algorithm	SHA.
IKE authentication method	Preshared keys.
IKE DH group identifier	Group 1.
IKE lifetime association	86,400 00 seconds (equals 24 hours).
IKE keepalive time for each peer (v2)	3,600 seconds (equals 1 hour).

Table 36-4 lists the default settings for IPsec parameters.

Table 36-4 Default IPsec Parameters

Parameters	Default
IPsec	Disabled.
Applying IPsec to the traffic.	Deny—allowing clear text.
IPsec PFS	Disabled.

Table 36-4 Default IPsec Parameters (continued)

Parameters	Default
IPsec global lifetime (traffic-volume)	450 Gigabytes.
IPsec global lifetime (time)	3,600 seconds (one hour).

CHAPTER 37

Configuring FC-SP and DHCHAP

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

This chapter includes the following sections:

About Fabric Authentication

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see Figure 37-1).

Figure 37-1 Switch and Host Authentication

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

DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, thus preventing unauthorized devices from accessing the switch.

Note The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see Chapter 3, "Obtaining and Installing Licenses").

To configure DHCHAP authentication using the local password database, follow these steps:

Step 1 Enable DHCHAP.

Step 2 Identify and configure the DHCHAP authentication modes.

Step 3 Configure the hash algorithm and DH group.

Step 4 Configure the DHCHAP password for the local switch and other switches in the fabric.

Step 5 Configure the DHCHAP timeout value for reauthentication.

Step 6 Verify the DHCHAP configuration.

This section includes the following topics:

DHCHAP Compatibility with Existing Cisco MDS Features

This sections identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- 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.
- High availability—DHCHAP authentication works transparently with existing HA features.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

Enabling DHCHAP

To enable DHCHAP for a Cisco MDS switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp enable`	Enables the DHCHAP in this switch.
	`switch(config)# no fcsp enable`	Disables (default) the DHCHAP in this switch.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to 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 such ports return error messages to the initiating switch.

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

Table 37-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Table 37-1 DHCHAP Authentication Status Between Two MDS Switches

Switch N DHCHAP Modes	Switch 1 DHCHAP Modes			
	on	**auto-active**	**auto-passive**	**off**
on	FC-SP authentication is performed.	FC-SP authentication is performed.	FC-SP authentication is performed.	Link is brought down.
auto-Active				FC-SP authentication is *not* performed.
auto-Passive			FC-SP authentication is *not* performed.	
off	Link is brought down.	FC-SP authentication is *not* performed.		

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc2/1-3` `switch(config-if)#`	Select a range of interfaces and enters the interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# fcsp on`	Sets the DHCHAP mode for the selected interfaces to be in the on state.
	`switch(config-if)# no fcsp on`	Reverts to the factory default of auto-passive for these three interfaces.
Step 4	`switch(config-if)# fcsp auto-active 0`	Changes the DHCHAP authentication mode for the selected interfaces to auto-active. Zero (0) indicates that the port does not perform reauthentication. **Note** The reauthorization interval configuration is the same as the default behavior.
	`switch(config-if)# fcsp auto-active 120`	Changes the DHCHAP authentication mode to auto-active for the selected interfaces and enables reauthentication every two hours (120 minutes) after the initial authentication.
	`switch(config-if)# fcsp auto-active`	Changes the DHCHAP authentication mode to auto-active for the selected interfaces. Reauthentication is disabled (default). **Note** The reauthorization interval configuration is the same as setting it to zero (0).

About the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.

Tip If you change the hash algorithm configuration, then change it globally for all switches in the fabric.

Caution RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp dhchap hash sha1`	Configures the use of only the SHA-1 hash algorithm.
	`switch(config)# fcsp dhchap hash MD5`	Configures the use of only the MD5 hash algorithm.
	`switch(config)# fcsp dhchap hash md5 sha1`	Defines the use of the default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.
	`switch(config)# no fcsp dhchap hash sha1`	Reverts to the factory default priority list of the MD5 hash algorithm followed by the SHA-1 hash algorithm.

About the DHCHAP Group Settings

All switches in the Cisco MDS Family 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.

Tip If you change the DH group configuration, change it globally for all switches in the fabric.

Configuring the DHCHAP Group Settings

To change the DH group settings, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp dhchap group 2 3 4`	Prioritizes the use of DH group 2, 3, and 4 in the configured order.
	`switch(config)# no fcsp dhchap group 0`	Reverts to the DHCHAP factory default order of 0, 4, 1, 2, and 3.

About the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric.
- Approach 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.
- Approach 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 approach requires considerable password maintenance by the user.

Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

Tip 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 Approach 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp dhchap password 0 mypassword`	Configures a clear text password for the local switch.
	`switch(config)# fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22`	Configures a clear text password for the local switch to be used for the device with the specified WWN.
	`switch(config)# no fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22`	Removes the clear text password for the local switch to be used for the device with the specified WWN.
	`switch(config)# fcsp dhchap password 7 sfsfdf`	Configures a password entered in an encrypted format for the local switch.
	`switch(config)# fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22`	Configures a password entered in an encrypted format for the local switch to be used for the device with the specified WWN.
	`switch(config)# no fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22`	Removes the password entered in an encrypted format for the local switch to be used for the device with the specified WWN.
	`switch(config)# fcsp dhchap password mypassword1`	Configures a clear text password for the local switch to be used with any connecting device.

About Password Configuration for Remote Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

Note The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

Configuring DHCHAP Passwords for Remote Devices

To locally configure the remote DHCHAP password for another switch in the fabric, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword`	Configures a password for another switch in the fabric that is identified by the switch WWN device name.
	`switch(config)# no fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword`	Removes the password entry for this switch from the local authentication database.
	`switch(config)# fcsp dhchap devicename 00:11:55:66:00:aa:bb:cc password 0 NewPassword`	Configures a clear text password for another switch in the fabric that is identified by the switch WWN device name.
	`switch(config)# fcsp dhchap devicename 00:11:22:33:55:aa:bb:cc password 7 asdflkjh`	Configures a password entered in an encrypted format for another switch in the fabric that is identified by the switch WWN device name.

About the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured on all switches in the fabric.

Configuring the DHCHAP Timeout Value

To configure the DHCHAP timeout value, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcsp timeout 60`	Configures the reauthentication timeout to be 60 seconds.
	`switch(config)# no fcsp timeout 60`	Reverts to the factory default of 30 seconds.

Configuring DHCHAP AAA Authentication

You can individually set authentication options. If authentication is not configured, local authentication is used by default.

To configure the AAA authentication , follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# aaa authentication dhchap default group TacacsServer1`	Enables DHCHAP to use the TACACS+ server group (in this example, TacacsServer1) for authentication.
	`switch(config)# aaa authentication dhchap default local`	Enables DHCHAP for local authentication.
	`switch(config)# aaa authentication dhchap default group RadiusServer1`	Enables DHCHAP to use the RADIUS server group (in this example, RadiusServer1) for authentication.

Displaying Protocol Security Information

Use the **show fcsp** commands to display configurations for the local database (see Example 37-1 through 37-6).

Example 37-1 Displays DHCHAP Configurations in FC Interfaces

```
switch# show fcsp interface fc1/9

fc1/9:
        fcsp authentication mode:SEC_MODE_ON
        Status: Successfully authenticated
```

Example 37-2 Displays DHCHAP Statistics for an FC Interface

```
switch# show fcsp interface fc1/9 statistics

fc1/9:
        fcsp authentication mode:SEC_MODE_ON
        Status: Successfully authenticated
        Statistics:
        FC-SP Authentication Succeeded:5
        FC-SP Authentication Failed:0
        FC-SP Authentication Bypassed:0
```

Example 37-3 Displays the FC-SP WWN of the Device Connected through a Specified Interface

```
switch# show fcsp interface fc 2/1 wwn

fc2/1:
        fcsp authentication mode:SEC_MODE_ON
        Status: Successfully authenticated
        Other device's WWN:20:00:00:e0:8b:0a:5d:e7
```

Example 37-4 Displays Hash Algorithm and DHCHAP Groups Configured for the Local Switch

```
switch# show fcsp dhchap
Supported Hash algorithms (in order of preference):
DHCHAP_HASH_MD5
DHCHAP_HASH_SHA_1

Supported Diffie Hellman group ids (in order of preference):
DHCHAP_GROUP_NULL
DHCHAP_GROUP_1536
DHCHAP_GROUP_1024
DHCHAP_GROUP_1280
DHCHAP_GROUP_2048
```

Example 37-5 Displays the DHCHAP Local Password Database

```
switch# show fcsp dhchap database
DHCHAP Local Password:
        Non-device specific password:**********
        Password for device with WWN:29:11:bb:cc:dd:33:11:22 is ******
        Password for device with WWN:30:11:bb:cc:dd:33:11:22 is *********

Other Devices' Passwords:
        Password for device with WWN:00:11:22:33:44:aa:bb:cc is ***********
```

Example 37-6 Displays the ASCII Representation of the Device WWN

```
switch# show fcsp asciiwwn 30:11:bb:cc:dd:33:11:22
Ascii representation of WWN to be used with AAA servers:Ox_3011bbccdd331122
```

Tip Use the ASCII representation of the device WWN (identified in bold in Example 37-6) to configure the switch information on RADIUS and TACACS+ servers.

Sample Configuration

This section provides the steps to configure the example illustrated in Figure 37-2.

Figure 37-2 Sample DHCHAP Authentication

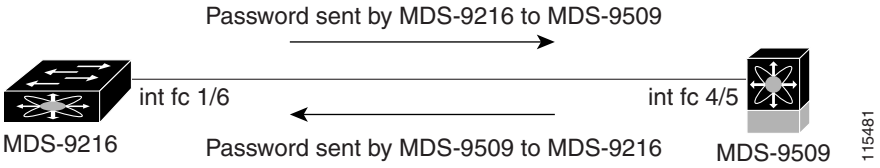

To configure the authentication setup shown in Figure 37-2, follow these steps:

Step 1 Obtain the device name of the MDS 9216 Switch in the fabric, The MDS 9216 Switch in the fabric is identified by the switch WWN.

```
MDS-9216# 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.

```
MDS-9216(config)# fcsp enable
```

Step 3 Configure a clear text password for this switch. This password will be used by the connecting device.

```
MDS-9216(config)# fcsp dhchap password rtp9216
```

Step 4 Configures a password for another switch in the fabric that is identified by the switch WWN device name.

```
MDS-9216(config)# fcsp dhchap devicename 20:00:00:05:30:00:38:5e password rtp9509
```

Step 5 Enable the DHCHAP mode for the required Fibre Channel interface.

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

```
MDS-9216(config)# interface fc 1/16
MDS-9216(config-if)# fcsp on
```

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

```
MDS-9216# show fcsp dhchap database
DHCHAP Local Password:
        Non-device specific password:*******
Other Devices' Passwords:
        Password for device with WWN:20:00:00:05:30:00:38:5e is *******
```

Step 7 Display the DHCHAP configuration in the Fibre Channel interface,

```
MDS-9216# show fcsp interface fc 1/6
fc1/6
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

Step 8 Repeat these steps on the connecting MDS 9509 Switch.

```
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 fc 4/5
Fc4/5
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

You have now enabled and configured DHCHAP authentication for the sample setup in Figure 37-2.

Default Settings

Table 37-2 lists the default settings for all fabric security features in any switch.

Table 37-2 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 38

Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

Note Port security is only supported for Fibre Channel ports.

This chapter includes the following sections:

About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family in the following ways:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the ENTERPRISE_PKG license (see Chapter 3, “Obtaining and Installing Licenses”).

This section includes the following topics:

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 Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. So, for example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.

Note If you enable auto-learning before activating port security, you cannot activate until auto-learning is disabled.

Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

By activating the port security feature, the following apply:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
  - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

Tip 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 issue a **no shutdown** CLI command to bring that port back online.

Port Security Configuration Guidelines

The steps to configure port security depend on which features you are using. Auto-learning works differently if you are using CFS distribution.

This section includes the following topics:

Configuring Port Security with Auto-Learning and CFS Distribution

To configure port security, using auto-learning and CFS distribution, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 38-5.

Step 2 Enable CFS distribution. See the "Enabling Distribution" section on page 38-12.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 38-5.

Step 4 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 38-13. At this point, all switches are activated, and auto-learning.

Step 5 Wait until all switches and all hosts are automatically learned.

Step 6 Disable auto-learn on each VSAN. See the"Disabling Auto-learning" section on page 38-8.

Step 7 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 38-13. At this point, 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. See the "Port Security Database Copy" section on page 38-17.

Step 9 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 38-13. This ensures that the configure database is the same on all switches in the fabric.

Step 10 Copy the running configuration to the startup configuration, using the fabric option. This saves the port security configure database to the startup configuration on all switches in the fabric.

Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 38-5.

Step 2 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 38-5.

Step 3 Wait until all switches and all hosts are automatically learned.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 38-8.

Step 5 Copy the active database to the configure database on each VSAN. See the "Port Security Database Copy" section on page 38-17.

Step 6 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Step 7 Repeat Step 1 through Step 6 for all switches in the fabric.

Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 38-5.

Step 2 Manually configure all port security entries into the configure database on each VSAN. See the "Configuring Port Security with Manual Database Configuration" section on page 38-4.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Disabling Auto-learning" section on page 38-8.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 38-8.

Step 5 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Step 6 Repeat Step 1 through Step 5 for all switches in the fabric.

Enabling Port Security

By default, the port security feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable port security, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-security enable`	Enables port security on that switch.
	`switch(config)# no port-security enable`	Disables (default) port security on that switch.

Port Security Activation

This section includes the following topics:

- Activating Port Security, page 38-5
- Database Activation Rejection, page 38-6
- Forcing Port Security Activation, page 38-6
- Database Reactivation, page 38-6

Activating Port Security

To activate port security, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security activate vsan 1`	Activates the port security database for the specified VSAN, and automatically enables auto-learning.
	`switch(config)# port-security activate vsan 1 no-auto-learn`	Activates the port security database for the specified VSAN, and disables auto-learning.
	`switch(config)# no port-security activate vsan 1`	Deactivates the port security database for the specified VSAN, and automatically disables auto-learning.

Note If required, you can disable auto-learning (see the "Disabling Auto-learning" section on page 38-8).

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 PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

Forcing Port Security Activation

If the port security activation request is rejected, you can force the activation.

Note An activation using the **force** option can log out existing devices if they violate the active database.

You can view missing or conflicting entries using the **port-security database diff active vsan** command in EXEC mode.

To forcefully activate the port security database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security activate vsan 1 force`	Forces the VSAN 1 port security database to activate despite conflicts.

Database Reactivation

Tip If auto-learning is enabled, you cannot activate the database, without the **force** option until you disable auto-learning.

To reactivate the port security database , follow these steps:

Step 1 Disable auto-learning.

Step 2 Copy the active database to the configured database.

Tip If the active database is empty, you cannot perform this step.

Step 3 Make the required changes to the configuration database.

Step 4 Activate the database.

To reactivate the port security database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no port-security auto-learn vsan 1`	Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point.
Step 3	`switch(config)# exit` `switch# port-security database copy vsan 1`	Copies from the active to the configured database.
Step 4	`switch# config t` `switch(config)# port-security activate vsan 1`	Activates the port security database for the specified VSAN, and automatically enables auto-learning.

Auto-learning

This section contains the following topics:

About Enabling Auto-learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.
- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).

Tip If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the **force** option.

Enabling Auto-learning

To enable auto-learning, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security auto-learn vsan 1`	Enables auto-learning so the switch can learn about any device that is allowed to access VSAN 1. These devices are logged in the port security active database.

Disabling Auto-learning

To disable auto-learning, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no port-security auto-learn vsan 1`	Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point.

Auto-learning Device Authorization

Table 38-1 summarizes the authorized connection conditions for device requests.

Table 38-1 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

Table 38-1 Authorized Auto-Learning Device Requests (continued)

Condition	Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization
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 fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 38-2 summarizes the port security authorization results for this active database. The conditions listed refer to the conditions from Table 38-1.

Table 38-2 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.

Table 38-2 Authorization Results for Scenario (continued)

Device Connection Request	Authorization	Condition	Reason
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.

Port Security Manual Configuration

To configure port security on any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Identify the WWN of the ports that need to be secured.

Step 2 Secure the fWWN to an authorized nWWN or pWWN.

Step 3 Activate the port security database.

Step 4 Verify your configuration.

This section includes the following topics:

About WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port is allowed to log in to SAN switch port Fx, then that Nx port can only log in through the specified Fx port..
- If an Nx port's nWWN is bound to an Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Adding Authorized Port Pairs

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

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

To add authorized port pairs for port security, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security database vsan 1` `switch(config-port-security)#`	Enters the port security database mode for the specified VSAN.
	`switch(config)# no port-security database vsan 1` `switch(config)#`	Deletes the port security configuration database from the specified VSAN.
Step 3	`switch(config-port-security)# swwn 20:01:33:11:00:2a:4a:66 interface port-channel 5`	Configures the specified sWWN to only log in through PortChannel 5.
	`switch(config-port-security)# any-wwn interface fc1/1 - fc1/8`	Configures any WWN to log in through the specified interfaces.
	`switch(config-port-security)# pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e`	Configures the specified pWWN to only log in through the specified fWWN.
	`switch(config-port-security)# no pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e`	Deletes the specified pWWN configured in the previous step.
	`switch(config-port-security)# nwwn 26:33:22:00:55:05:3d:4c fwwn 20:81:00:44:22:00:4a:9e`	Configures the specified nWWN to log in through the specified fWWN.
	`switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66`	Configures the specified pWWN to log in through any port in the fabric.
	`switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80`	Configures the specified pWWN to log in through any interface in the specified switch.
	`switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80 interface fc3/1`	Configures the specified pWWN to log in through the specified interface in the specified switch
	`switch(config-port-security)# any-wwn interface fc3/1`	Configures any WWN to log in through the specified interface in any switch.
	`switch(config-port-security)# no any-wwn interface fc2/1`	Deletes the wildcard configured in the previous step.

Port Security Configuration Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric (see Chapter 6, “Using the CFS Infrastructure”).

This section contains the following topics:

Enabling Distribution

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

Note Port Activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration. See the "Activation and Auto-learning Configuration Distribution" section on page 38-13.

For example, if you activate port security, follow up by disabling auto-learning, and finally commit the changes in the pending database, then the net result of your actions is the same as issuing a **port-security activate vsan** *vsan-id* **no-auto-learn** command.

Tip In this case, we recommend that you perform a commit at the end of each operation: After you activate port security and after you enable auto learning.

To enable the port security distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security distribute`	Enables distribution.
	`switch(config)# no port-security distribute`	Disables 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

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the port security configuration changes for the specified VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security commit vsan 3`	Commits the port security changes in the specified VSAN.

Discarding the Changes

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

To discard the port security configuration changes for the specified VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# port-security abort vsan 5`	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 merely 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, then the activation and auto-learning changes are consolidated and the behavior may change (see Table 38-3).

Table 38-3 Scenarios for Activation and Auto-learning Configurations in Distributed Mode

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C[1], D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. A new entry E is added to the configuration database.	configuration database = {A,B, E} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. You disable learning.	configuration database = {A,B} active database = {A,B, C, D}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled + learning to be disabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled. pending database = empty

1. The * (asterisk) indicates learned entries.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the "CFS Merge Support" section on page 6-8 for detailed concepts.

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 2K.

Caution If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Database Interaction

Table 38-4 lists the differences and interaction between the active and configuration databases.

Table 38-4 Active and Configuration Port Security Databases

Active Database	Configuration Database
Read-only.	Read-write.
Saving the configuration only saves the activated entries. Learned entries are not saved.	Saving the configuration saves all the entries in the configuration database.
Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.	Once activated, the configuration database can be modified without any effect on the active database.
You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.	You can overwrite the configuration database with the active database.

Note You can overwrite the configuration database with the active database using the **port-security database copy vsan** command. The **port-security database diff active vsa**n command in EXEC mode lists the differences between the active database and the configuration database.

This section includes the following topics:

Database Scenarios

Figure 38-1 depicts various scenarios to depict the active database and the configuration database status based on port security configurations.

Figure 38-1 Port Security Database Scenarios

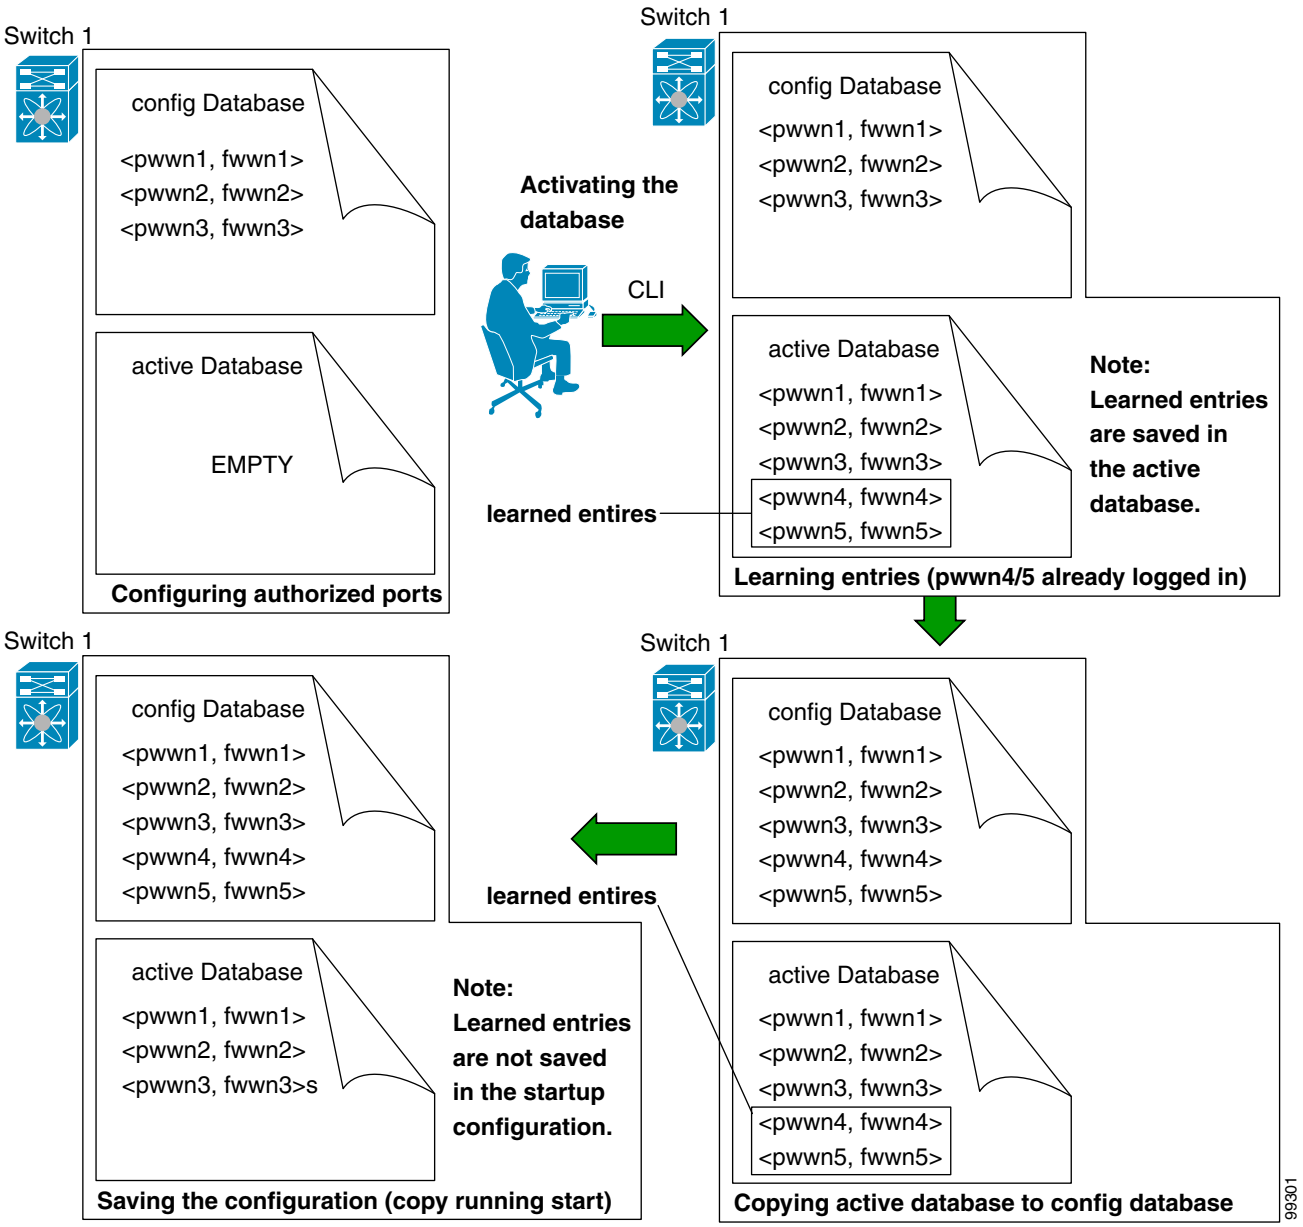

Port Security Database Copy

Tip We recommend that you issue the **port-security database copy vsan** command after disabling auto-learning. This action will ensure that the configuration database is in sync with the active database. If distribution is enabled, this command creates a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

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

```
switch# port-security database diff config vsan 1
```

Port Security Database Deletion

Tip If the distribution is enabled, the deletion creates a copy of the database. An explicit **port-security commit** command is required to actually delete the database.

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

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

Port Security Database Cleanup

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

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

Use the **clear port-security database auto-learn interface** command to clear any learned entries in the active database for a specified interface within a VSAN.

```
switch# clear port-security database auto-learn interface fc1/1 vsan 1
```

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

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

Note The **clear port-security database auto-learn** and **clear port-security statistics** commands are only relevant to the local switch and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

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

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

Displaying Port Security Configuration

The **show port-security database** commands display the configured port security information (see Examples 38-1 to 38-11).

Example 38-1 Displays the Contents of the Port Security Configuration Database

```
switch# show port-security database
-------------------------------------------------------------------------------------
VSAN    Logging-in Entity               Logging-in Point       (Interface)
-------------------------------------------------------------------------------------
1       21:00:00:e0:8b:06:d9:1d(pwwn)  20:0d:00:05:30:00:95:de(fc1/13)
1       50:06:04:82:bc:01:c3:84(pwwn)  20:0c:00:05:30:00:95:de(fc1/12)
2       20:00:00:05:30:00:95:df(swwn)  20:0c:00:05:30:00:95:de(port-channel 128)
3       20:00:00:05:30:00:95:de(swwn)  20:01:00:05:30:00:95:de(fc1/1)
[Total 4 entries]
```

You can optionally specify a fWWN and a VSAN, or an interface and a VSAN in the **show port-security** command to view the output of the activated port security (see Example 38-2).

Example 38-2 Displays the Port Security Configuration Database in VSAN 1

```
switch# show port-security database vsan 1
-------------------------------------------------------------------------------
Vsan   Logging-in Entity              Logging-in Point        (Interface)
-------------------------------------------------------------------------------
1                  *                  20:85:00:44:22:00:4a:9e (fc3/5)
1      20:11:00:33:11:00:2a:4a(pwwn)  20:81:00:44:22:00:4a:9e (fc3/1)
[Total 2 entries]
```

Example 38-3 Displays the Activated Database

```
switch# show port-security database active
-------------------------------------------------------------------------------------------
VSAN    Logging-in Entity               Logging-in Point     (Interface)         Learnt
-------------------------------------------------------------------------------------------
1       21:00:00:e0:8b:06:d9:1d(pwwn)   20:0d:00:05:30:00:95:de(fc1/13)           Yes
1       50:06:04:82:bc:01:c3:84(pwwn)   20:0c:00:05:30:00:95:de(fc1/12)           Yes
2       20:00:00:05:30:00:95:df(swwn)   20:0c:00:05:30:00:95:de(port-channel 128) Yes
3       20:00:00:05:30:00:95:de(swwn)   20:01:00:05:30:00:95:de(fc1/1)
[Total 4 entries]
```

Example 38-4 Displays the Contents of the Temporary Configuration Database

```
switch# show port-security pending vsan 1
Session Context for VSAN 1
--------------------------
Activation Status: Active
Auto Learn Status: On
Force activate: No
Config db modified: Yes
Activation done: Yes
Session owner: admin(2)
Session database:
--------------------------------------------------------------------------------
VSAN Logging-in Entity Logging-in Point (Interface)
--------------------------------------------------------------------------------
1 20:11:00:33:22:00:2a:4a(pwwn) 20:41:00:05:30:00:4a:1e(fc2/1)
[Total 1 entries]
```

Example 38-5 Displays the Difference Between the Temporary Configuration Database and the Configuration Database

```
switch# show port-security pending-diff vsan 1
Session Diff for VSAN: 1
------------------------
Database will be activated
Learning will be turned ON
Database Diff:
+pwwn 20:11:00:33:22:00:2a:4a fwwn 20:41:00:05:30:00:4a:1e
```

The 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 (see Examples 38-6 to 38-8).

Example 38-6 Displays the Wildcard fWWN Port Security in VSAN 1

```
switch# show port-security database fwwn 20:85:00:44:22:00:4a:9e vsan 1
Any port can login thru' this fwwn
```

Example 38-7 Displays the Configured fWWN Port Security in VSAN 1

```
switch# show port-security database fwwn 20:01:00:05:30:00:95:de vsan 1
20:00:00:0c:88:00:4a:e2(swwn)
```

Example 38-8 Displays the Interface Port Information in VSAN 2

```
switch# show port-security database interface fc 1/1 vsan 2
20:00:00:0c:88:00:4a:e2(swwn)
```

The port security statistics are constantly updated and available at any time (see Example 38-9).

Example 38-9 Displays the Port Security Statistics

```
switch# show port-security statistics
Statistics For VSAN: 1
-----------------------
Number of pWWN permit: 2
Number of nWWN permit: 2
Number of sWWN permit: 2
Number of pWWN deny  : 0
Number of nWWN deny  : 0
Number of sWWN deny  : 0

Total Logins permitted  : 4
Total Logins denied     : 0
Statistics For VSAN: 2
-----------------------
Number of pWWN permit: 0
Number of nWWN permit: 0
Number of sWWN permit: 2
Number of pWWN deny  : 0
Number of nWWN deny  : 0
Number of sWWN deny  : 0
...
```

To verify the status of the active database and the auto-learning configuration, use the **show port-security status** command (see Example 38-10).

Example 38-10 Displays the Port Security Status

```
switch# show port-security status
Fabric Distribution Enabled
VSAN 1 :No Active database, learning is disabled, Session Lock Taken
VSAN 2 :No Active database, learning is disabled, Session Lock Taken
...
```

The **show port-security** command displays the previous 100 violations by default (see Example 38-11).

Example 38-11 Displays the Violations in the Port Security Database

```
switch# show port-security violations

-------------------------------------------------------------------------------------------
VSAN    Interface       Logging-in Entity               Last-Time          [Repeat count]
-------------------------------------------------------------------------------------------
1       fc1/13          21:00:00:e0:8b:06:d9:1d(pwwn)   Jul  9 08:32:20 2003    [20]
                        20:00:00:e0:8b:06:d9:1d(nwwn)
1       fc1/12          50:06:04:82:bc:01:c3:84(pwwn)   Jul  9 08:32:20 2003    [1]
                        50:06:04:82:bc:01:c3:84(nwwn)
2       port-channel 1  20:00:00:05:30:00:95:de(swwn)   Jul  9 08:32:40 2003    [1]
[Total 2 entries]
```

The **show port-security** command issued with the **last** *number* option displays only the specified number of entries that appear first.

Default Settings

Table 38-5 lists the default settings for all port security features in any switch.

Table 38-5 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 39

Configuring Fabric Binding

This chapter describes the fabric binding feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

About Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

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.

This section has the following topics:

Licensing Requirements

Fabric binding requires that you install either the MAINFRAME_PKG license or the ENTERPRISE_PKG license on your switch.

See Chapter 3, “Obtaining and Installing Licenses,” for more information on license feature support and installation.

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. Table 39-1 compares the two features.

Table 39-1 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 ports. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for xE ports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
  - E port security binding check on port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. In a FICON VSAN, the fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database. In a Fibre Channel VSAN, only the sWWN is required; the domain ID is optional.

Note All switches in a Fibre Channel VSAN using fabric binding must be running Cisco MDS SAN-OS Release 3.0(1) or later.

Fabric Binding Configuration

To configure fabric binding in each switch in the fabric, follow these steps.

Step 1 Enable the fabric configuration feature.

Step 2 Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.

Step 3 Activate the fabric binding database.

Step 4 Copy the fabric binding active database to the fabric binding config database.

Step 5 Save the fabric binding configuration.

Step 6 Verify the fabric binding configuration.

Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fabric-binding enable`	Enables fabric binding on that switch.
	`switch(config)# no fabric-binding enable`	Disables (default) fabric binding on that switch.

View the status of the fabric binding feature of a fabric binding-enabled switch by issuing the **show fabric-binding status** command.

```
switch# show fabric-binding status
VSAN 1:Activated database
VSAN 4:No Active database
```

Configuring Switch WWN List

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.

The persistent domain ID can be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured and the end devices reject a domain ID change in all switches in the fabric. Domain ID authorization is not required in Fibre Channel VSANs.

To configure a list of sWWNs and domain IDs for a FICON VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fabric-binding database vsan 5` `switch(config-fabric-binding)#`	Enters the fabric binding submode for the specified VSAN.
	`switch(config)# no fabric-binding database vsan 5`	Deletes the fabric binding database for the specified VSAN.
Step 3	`switch(config-fabric-binding)# swwn 21:00:05:30:23:11:11:11 domain 102`	Adds the sWWN and domain ID of a switch to the configured database list.
	`switch(config-fabric-binding)# swwn 21:00:05:30:23:1a:11:03 domain 101`	Adds the sWWN and domain ID of another switch to the configured database list.
	`switch(config-fabric-binding)# no swwn 21:00:15:30:23:1a:11:03 domain 101`	Deletes the sWWN and domain ID of a switch from the configured database list.
Step 4	`switch(config-fabric-binding)# exit` `switch(config)#`	Exits the fabric binding submode.

To configure a list of sWWNs and optional domain IDs for a Fibre Channel VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fabric-binding database vsan 10` `switch(config-fabric-binding)#`	Enters the fabric binding submode for the specified VSAN.
	`switch(config)# no fabric-binding database vsan 10`	Deletes the fabric binding database for the specified VSAN.
Step 3	`switch(config-fabric-binding)# swwn 21:00:05:30:23:11:11:11`	Adds the sWWN of a switch for all domains to the configured database list.
	`switch(config-fabric-binding)# no swwn 21:00:05:30:23:11:11:11`	Deletes the sWWN of a switch for all domains from the configured database list.
	`switch(config-fabric-binding)# swwn 21:00:05:30:23:1a:11:03 domain 101`	Adds the sWWN of another switch for a specific domain ID to the configured database list.
	`switch(config-fabric-binding)# no swwn 21:00:15:30:23:1a:11:03 domain 101`	Deletes the sWWN and domain ID of a switch from the configured database list.
Step 4	`switch(config-fabric-binding)# exit` `switch(config)#`	Exits the fabric binding submode.

Fabric Binding Activation

The fabric binding feature maintains a configuration database (config-database) and an active database. The config-database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config- database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the fabric binding database on the switch if entries existing in the configured database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config-database. You can choose to forcefully override these situations.

Note After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

To activate the fabric binding feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fabric-binding activate vsan 10`	Activates the fabric binding database for the specified VSAN.
	`switch(config)# no fabric-binding activate vsan 10`	Deactivates the fabric binding database for the specified VSAN.

Forcing Fabric Binding Activation

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

To forcefully activate the fabric binding database, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fabric-binding activate vsan 3 force`	Activates the fabric binding database for the specified VSAN forcefully—even if the configuration is not acceptable.
	`switch(config)# no fabric-binding activate vsan 3 force`	Reverts to the previously configured state or to the factory default (if no state is configured).

Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config database is saved to the running configuration.

Caution You cannot disable fabric binding in a FICON-enabled VSAN.

- Use the **fabric-binding database copy vsan** command to copy from the active database to the config database. If the configured database is empty, this command is not accepted.

  ```
  switch# fabric-binding database copy vsan 1
  ```

- Use the **fabric-binding database diff active vsan** command to view the differences between the active database and the config database. This command can be used when resolving conflicts.

  ```
  switch# fabric-binding database diff active vsan 1
  ```

- Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the config database and the active database.

  ```
  switch# fabric-binding database diff config vsan 1
  ```

- Use the **copy running-config startup-config** command to save the running configuration to the startup configuration so that the fabric binding config database is available after a reboot.

  ```
  switch# copy running-config startup-config
  ```

Clearing the Fabric Binding Statistics

Use the **clear fabric-binding statistics** command to clear all existing statistics from the fabric binding database for a specified VSAN.

```
switch# clear fabric-binding statistics vsan 1
```

Deleting the Fabric Binding Database

Use the **no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no fabric-binding database vsan 10
```

Verifying Fabric Binding Configurations

Use the **show** commands to display all fabric binding information configured on this switch (see Examples 39-1 to 39-9).

Example 39-1 Displays Configured Fabric Binding Database Information

```
switch# show fabric-binding database
-------------------------------------------------
Vsan    Logging-in Switch WWN     Domain-id
-------------------------------------------------
1       21:00:05:30:23:11:11:11   0x66(102)
1       21:00:05:30:23:1a:11:03    0x19(25)
1       20:00:00:05:30:00:2a:1e   0xea(234) [Local]
4       21:00:05:30:23:11:11:11         Any
4       21:00:05:30:23:1a:11:03         Any
4       20:00:00:05:30:00:2a:1e   0xea(234) [Local]
61      21:00:05:30:23:1a:11:03    0x19(25)
61      21:00:05:30:23:11:11:11   0x66(102)
61      20:00:00:05:30:00:2a:1e   0xea(234) [Local]
[Total 7 entries]
```

Example 39-2 Displays Active Fabric Binding Information

```
switch# show fabric-binding database active
-------------------------------------------------
Vsan    Logging-in Switch WWN     Domain-id
-------------------------------------------------
1       21:00:05:30:23:11:11:11   0x66(102)
1       21:00:05:30:23:1a:11:03    0x19(25)
1       20:00:00:05:30:00:2a:1e   0xea(234) [Local]
61      21:00:05:30:23:1a:11:03    0x19(25)
61      21:00:05:30:23:11:11:11   0x66(102)
61      20:00:00:05:30:00:2a:1e   0xef(239) [Local]
```

Example 39-3 Displays Configured VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database vsan 4
-------------------------------------------------
Vsan    Logging-in Switch WWN     Domain-id
-------------------------------------------------
4       21:00:05:30:23:11:11:11          Any
4       21:00:05:30:23:1a:11:03          Any
4       20:00:00:05:30:00:2a:1e   0xea(234) [Local]
[Total 2 entries]
```

Example 39-4 Displays Active VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database active vsan 61
-------------------------------------------------
Vsan    Logging-in Switch WWN     Domain-id
-------------------------------------------------
61      21:00:05:30:23:1a:11:03    0x19(25)
61      21:00:05:30:23:11:11:11   0x66(102)
61      20:00:00:05:30:00:2a:1e   0xef(239) [Local]
[Total 3 entries]
```

Example 39-5 Displays Fabric Binding Statistics

```
switch# show fabric-binding statistics
Statistics For VSAN: 1
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 4
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 61
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 345
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 346
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 347
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 348
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 789
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted  : 0
Total Logins denied     : 0
Statistics For VSAN: 790
-----------------------
Number of sWWN permit: 0
Number of sWWN deny  : 0
```

```
Total Logins permitted  : 0
Total Logins denied     : 0
```

Example 39-6 Displays Fabric Binding Status for Each VSAN

```
switch# show fabric-binding status
VSAN 1 :Activated database
VSAN 4 :No Active database
VSAN 61 :Activated database
VSAN 345 :No Active database
VSAN 346 :No Active database
VSAN 347 :No Active database
VSAN 348 :No Active database
VSAN 789 :No Active database
VSAN 790 :No Active database
```

Example 39-7 Displays Fabric Binding Violations

```
switch# show fabric-binding violations
--------------------------------------------------------------------------------
VSAN Switch WWN [domain]     Last-Time             [Repeat count] Reason
--------------------------------------------------------------------------------
2    20:00:00:05:30:00:4a:1e [0xeb] Nov 25 05:46:14 2003   [2]    Domain mismatch
3    20:00:00:05:30:00:4a:1e [*] Nov 25 05:44:58 2003      [2]    sWWN not found
4    20:00:00:05:30:00:4a:1e [*] Nov 25 05:46:25 2003      [1]    Database mismatch
```

Note In VSAN 3 the sWWN itself was not found in the list. In VSAN 2, the sWWN was found in the list, but has a domain ID mismatch.

Example 39-8 Displays EFMD Statistics

```
switch# show fabric-binding efmd statistics

EFMD Protocol Statistics for VSAN 1
---------------------------------------
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

EFMD Protocol Statistics for VSAN 4
---------------------------------------
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

EFMD Protocol Statistics for VSAN 61
---------------------------------------
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0
```

Example 39-9 Displays EFMD Statistics for a Specified VSAN

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

EFMD Protocol Statistics for VSAN 4
---------------------------------------
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0
```

Default Settings

Table 39-2 lists the default settings for the fabric binding feature.

Table 39-2 Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled.

PART 6

IP Services

CHAPTER 40

Configuring FCIP

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using o pen-standard, IP-based technology. The switch can connect separated SAN islands using Fibre Channel over IP (FCIP).

Note FCIP is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Di rectors.

The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switc h in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

Note For information on configuring Gigabit Ethernet interfaces, see Chapter 45, “Configuring IPv4 for Gi gabit Ethernet Interfaces.”

This chapter includes the following sections:

About FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distri buted Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). See Figure 40-1.

Figure 40-1 Fibre Channel SANs Connected by FCIP

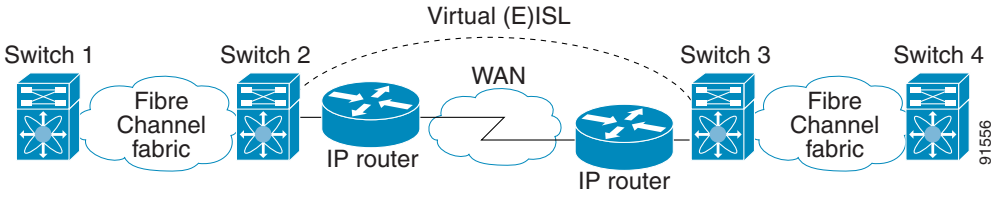

FCIP uses TCP as a network layer transport.

Note For more information about FCIP protocols, refer to the IETF standards for IP storage at http://www.ietf.org. Also refer to Fibre Channel standards for switch backbone connection at http://www.t11.org (see FC-BB-2).

This section includes the following topics:

- FCIP Concepts, page 40-2
- FCIP High-Availability Solutions, page 40-4
- Ethernet PortChannels and Fibre Channel PortChannels, page 40-7

FCIP Concepts

To configure IPS modules or MPS-14/2 modules for FCIP, you should have a basic understanding of the following concepts:

- FCIP and VE Ports, page 40-2
- FCIP Links, page 40-3
- FCIP Profiles, page 40-4
- FCIP Interfaces, page 40-4

FCIP and VE Ports

Figure 40-2 describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see Figure 40-2).

Figure 40-2 FCIP Links and Virtual ISLs

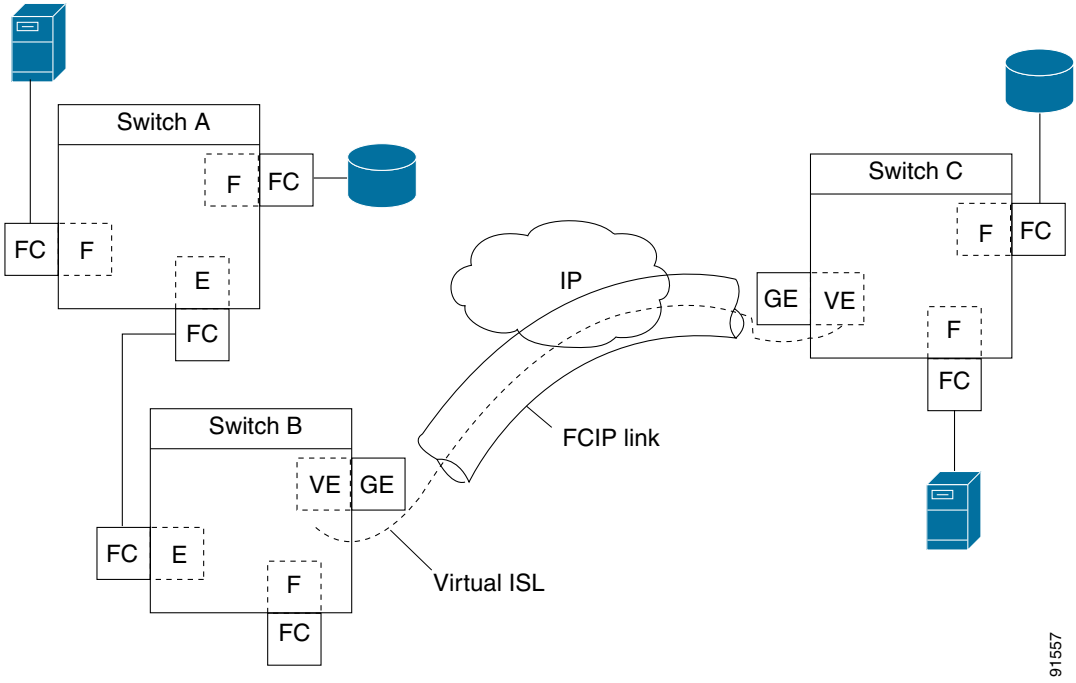

See the "E Port" section on page 12-4.

FCIP Links

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link:

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch proto col frames (all Class F). This arrangement provides low latency for all control frames.

To enable FCIP on the IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a no rmal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based o n the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-swi tch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

FCIP Profiles

The FCIP profile contains information about the local IP address and TCP parameters. The profile defi nes the following information:

- The local connection points (IP address and TCP port number)
- The behavior of the underlying TCP connections for all FCIP links that use this profile

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see Figure 40-3).

Figure 40-3 FCIP Profile and FCIP Links

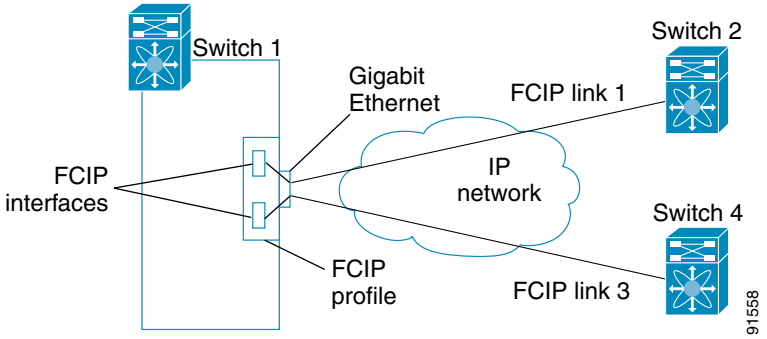

FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E po rt parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TC P connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

FCIP High-Availability Solutions

The following high-availability solutions are available for FCIP configurations:

- Fibre Channel PortChannels, page 40-5
- FSPF, page 40-5
- VRRP, page 40-6
- Ethernet PortChannels, page 40-6

Fibre Channel PortChannels

Figure 40-4 provides an example of a PortChannel-based load-balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 40-4 PortChannel-Based Load Balancing

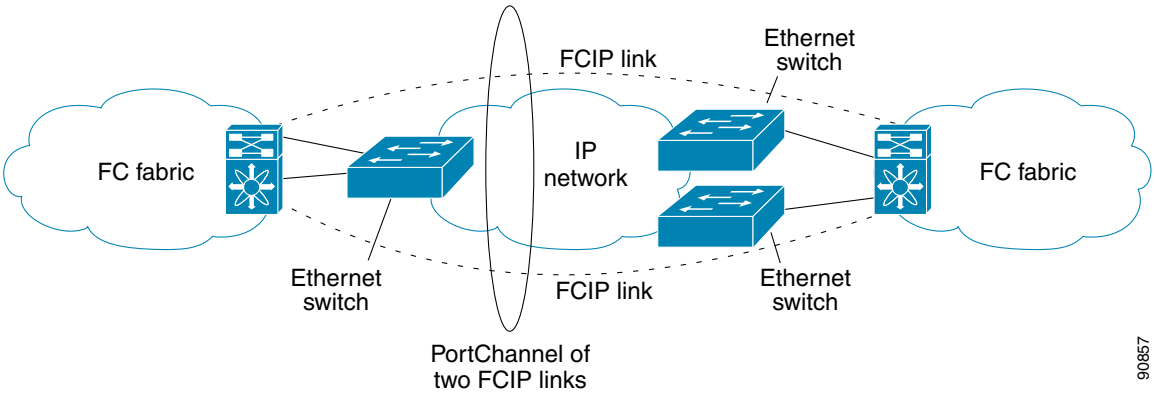

The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.

FSPF

Figure 40-5 displays a FPSF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 40-5 FSPF-Based Load Balancing

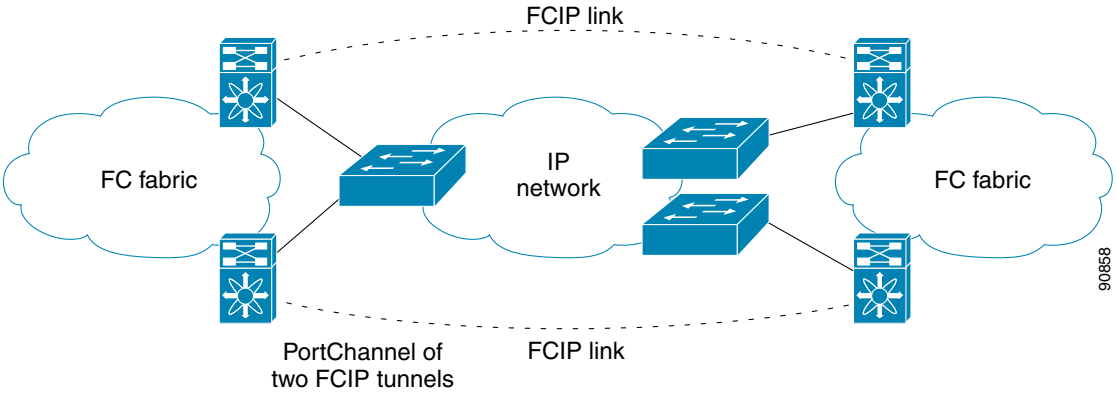

The following characteristics set FSPF solutions apart from other solutions:

- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP link.

VRRP

Figure 40-6 displays a Virtual Router Redundancy Protocol (VRRP)-based high availability FCIP confi guration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 40-6 VRRP-Based High Availability

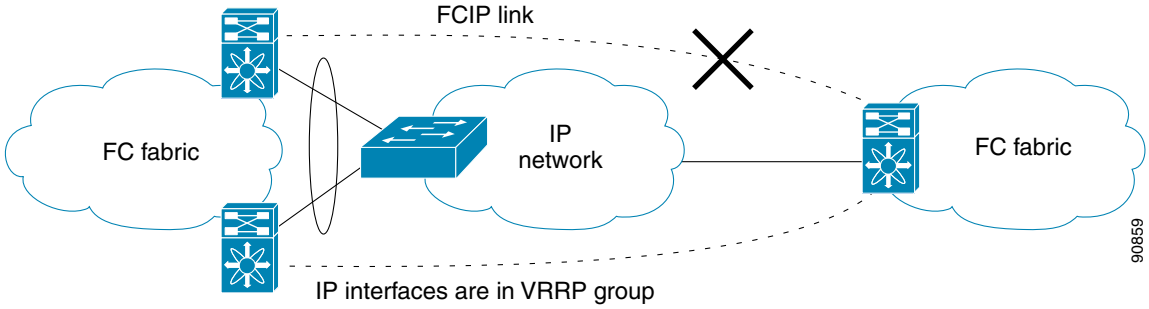

The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Ethernet PortChannels

Figure 40-7 displays an Ethernet PortChannel-based high- availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 40-7 Ethernet PortChannel-Based High Availability

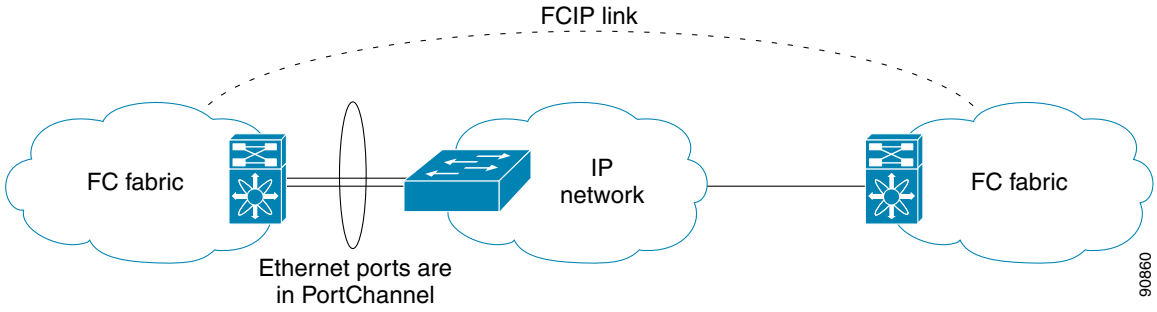

The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet li nk.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer link redundancy between the Cisco MDS 9000 Family switch's Gigabit Ethernet ports and the connecting Ethernet switch. On the other hand, Fibre Channel PortChannels also offer (E)ISL link redundancy between Fibre Channel switches. FCIP is an (E)ISL link and is only appli cable for a Fibre Channel PortChannel. Beneath the FCIP level, an FCIP link can run on top of an Ethernet PortChannel or just on one Gigabit Ethernet port. This link is totally transparent to the Fibre C hannel layer.

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet PortChannel (see the "Configuring Gigabit Ethernet High Availability" s ection on page 44-5). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortC hannel (to which FCIP link can be a part of) does not have a restriction on which (E)ISL links can be c ombined in a Fibre Channel PortChannel as long as it passes the compatibility check (see the "Compati bility Check" section on page 16-11). The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see Figure 40-8).

Figure 40-8 PortChannels at the Fibre Channel and Ethernet Levels

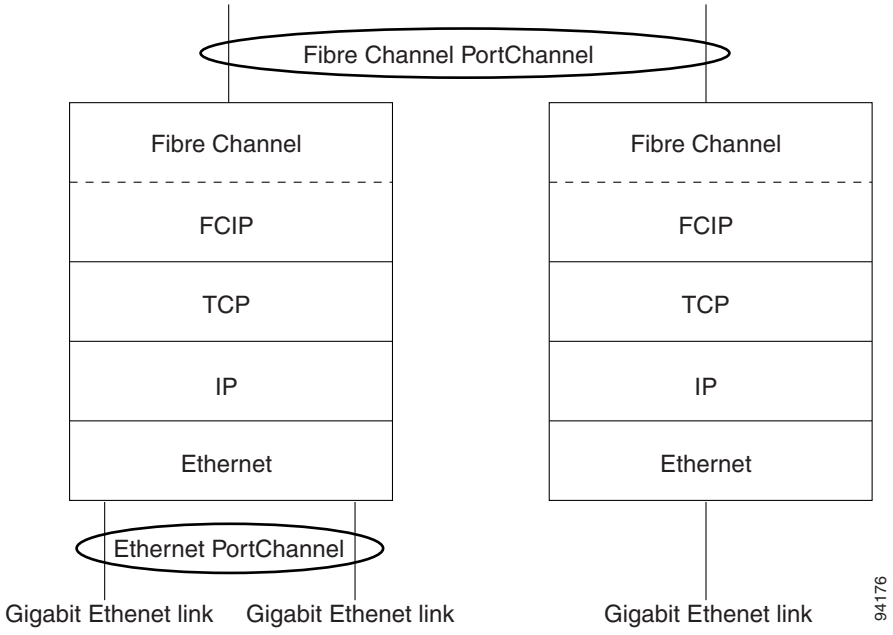

To configure Fibre Channel PortChannels, see Chapter 16, "Configuring PortChannels." To configure Ethernet PortChannels, see the "Configuring Gigabit Ethernet High Availability" section on page 44-5.

Configuring FCIP

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

Enabling FCIP

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

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

To use the FCIP feature, you need to obtain the SAN extension over IP package license (SAN_EXTN_OVER_IP or SAN_EXTN_OVER_IP_IPS4) (see Chapter 3, "Obtaining and Installing Li censes").

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

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcip enable`	Enables FCIP on that switch.
	`switch(config)# no fcip enable`	Disables (default) FCIP on that switch.

Note If FICON is enabled/FICON VSAN is present on both the switches, the Figure 40-15 is dis played, otherwise Figure 40-14 is displayed.

Basic FCIP Configuration

To configure an FCIP link, follow these steps on both switches:

Step 1 Configure the Gigabit Ethernet interface.

See the Chapter 45, "Configuring IPv4 for Gigabit Ethernet Interfaces."

Step 2 Create an FCIP profile, and then assign the Gigabit Ethernet interface's IP address to the profile.

Step 3 Create an FCIP interface, and then assign the profile to the interface.

Step 4 Configure the peer IP address for the FCIP interface.

Step 5 Enable the interface.

Creating FCIP Profiles

You must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile. You can assign IPv4 or IPv6 addresses to the interfaces. Figure 40-9 shows an example configuration.

Figure 40-9 Assigning Profiles to Each Gigabit Ethernet Interface

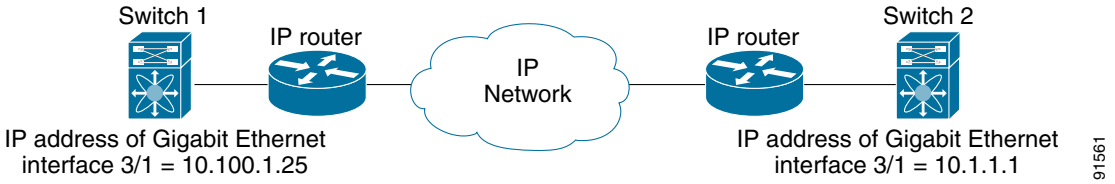

To create an FCIP profile in switch 1 in Figure 40-9, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch1(config)#`	Enters configuration mode.
Step 2	`switch1(config)# fcip profile 10` `switch1(config-profile)#`	Creates a profile for the FCIP connection. The valid range is from 1 to 255.
Step 3	`switch1(config-profile)# ip address 10.100.1.25`	Associates the profile (10) with the local IPv4 address of the Gigabit Ethernet interface (3/1).

To assign FCIP profile in switch 2 in Figure 40-9, follow these steps:

	Command	Purpose
Step 1	`switch2# config terminal` `switch2(config)#`	Enters configuration mode.
Step 2	`switch2(config)# fcip profile 20` `switch2(config-profile)#`	Creates a profile for the FCIP connection.
Step 3	`switch2(config-profile)# ip address 10.1.1.1`	Associates the profile (20) with the local IPv4 address of the Gigabit Ethernet interface.

Displaying FCIP Profile Information

Example 40-1 Displays FCIP Profiles

```
switch# show fcip profile

-------------------------------------------------------------------------------
ProfileId       Ipaddr          TcpPort
-------------------------------------------------------------------------------
1               10.10.100.150   3225
2               10.10.100.150   3226
40              40.1.1.2        3225
100             100.1.1.2       3225
200             200.1.1.2       3225
```

Example 40-2 Displays the Specified FCIP Profile Information

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

Creating FCIP Links

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules o r MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (s ee Figure 40-10).

Figure 40-10 Assigning Profiles to Each Gigabit Ethernet Interface

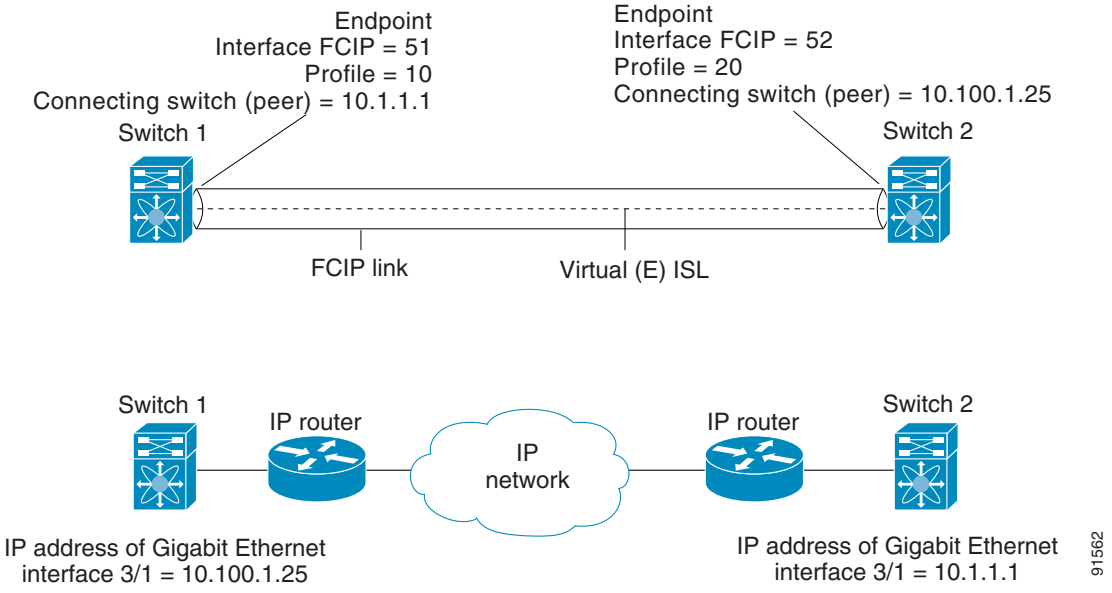

To create FCIP link endpoint in switch 1, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch1(config-if)# use-profile 10`	Assigns the profile (10) to the FCIP interface.

	Command	Purpose
Step 4	`switch1(config-if)# peer-info ipaddr 10.1.1.1`	Assigns the peer IPv4 address information (10.1.1.1 for switch 2) to the FCIP interface.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

To create an FCIP link endpoint in switch 2, follow these steps:

	Command	Purpose
Step 1	`switch2# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch2(config)# interface fcip 52` `switch2(config-if)#`	Creates an FCIP interface (52).
Step 3	`switch2(config-if)# use-profile 20`	Binds the profile (20) to the FCIP interface.
Step 4	`switch2(config-if)# peer-info ip addres s 10.100.1.25`	Assigns the peer IPv4 address information (10.100.1.25 for switch 1) to the FCIP interface.
Step 5	`switch1(config-if)# no shutdown`	Enables the interface.

Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the lo cal IP address and the local port, you can specify other TCP parameters as part of the FCIP profile confi guration.

- Configuring TCP Listener Ports, page 40-11
- Configuring TCP Parameters, page 40-12
- Displaying FCIP Profile Configuration Information, page 40-16

FCIP configuration options can be accessed from the `switch(config-profile)#` submode prompt.

Configuring TCP Listener Ports

To configure TCP listener ports, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcip profile 20` `switch(config-profile)#`	Creates the profile (if it does not already exist) and enters profi le configuration submode. The valid range is from 1 to 255.

The default TCP port for FCIP is 3225. You can change this port using the **port** command.

To change the default FCIP port number (3225), follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# port 5000`	Associates the profile with the local port number (5000).
	`switch(config-profile)# no port`	Reverts to the default 3225 port.

Configuring TCP Parameters

You can control TCP behavior in a switch by configuring the following TCP parameters.

Note When FCIP is sent over a WAN link, the default TCP settings may not be appropriate. In such cases, we recommend that you tune the FCIP WAN link by modifying the TCP parameters (specifically bandwi dth, round-trip times, and CWM burst size).

This section includes the following topics:

Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds (msec).

To configure the minimum retransmit time, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp mi n-retransmit-time 500`	Specifies the minimum TCP retransmit time for the TCP co nnection to be 500 msec. The default is 200 msec and the range is from 200 to 5000 msec.
	`switch(config-profile)# no tcp min-retransmit-time 500`	Reverts the minimum TCP retransmit time to the factory default of 200 msec.

Keepalive Timeout

You can configure the interval that the TCP connection uses to verify that the FCIP link is functioning. This ensures that an FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to tune the time taken to detect FCIP link failures.

You can configure the first interval during which the connection is idle (the default is 60 seconds). When the connection is idle for the configured interval, eight keepalive probes are sent at 1-second intervals. If no response is received for these eight probes and the connection remains idle throughout, that FCIP link is automatically closed.

Note Only the first interval (during which the connection is idle) can be changed.

To configure the first keepalive timeout interval, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp keepalive-timeout 120`	Specifies the keepalive timeout interval for the TCP connection in seconds (120). The range is from 1 to 7200 seconds.
	`switch(config-profile)# no tcp keepalive-timeout 120`	Reverts the keepalive timeout interval to the default 60 seconds.

Maximum Retransmissions

You can specify the maximum number of times a packet is retransmitted before TCP decides to close the connection.

To configure maximum retransmissions, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp max-retransmissions 6`	Specifies the maximum number of retransmissions (6). The range is from 1 to 8 retransmissions.
	`switch(config-profile)# no tcp max-retransmissions 6`	Reverts to the default of 4 retransmissions.

Path MTUs

Path MTU (PMTU) is the minimum MTU on the IP network between the two endpoints of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a timeout of 3600 seconds. If TCP reduces the size of the maximum segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

To configure PMTU, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp pmtu-enable`	Disables PMTU discovery.
	`switch(config-profile)# tcp pmtu-enable`	Enables (default) PMTU discovery with the default value of 3600 seconds.
	`switch(config-profile)# tcp pmtu-enable reset-timeout 90`	Specifies the PMTU reset timeout to 90 seconds. The range is 60 to 3600 seconds.
	`switch(config-profile)# no tcp pmtu-enable reset-timeout 600`	Leaves PMTU discovery enabled but reverts the timeout to the default of 3600 seconds.

Selective Acknowledgments

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

To configure SACK, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp sack-enable`	Disables SACK.
	`switch(config-profile)# tcp sack-enable`	Enables SACK (default).

Window Management

The optimal TCP window size is automatically calculated using the maximum bandwidth parameter, the minimum available bandwidth parameter, and the dynamically measured round trip time (RTT).

Note The configured **round-trip-time** parameter determines the window scaling factor of the TCP connectio n. This parameter is only an approximation. The measured RTT value overrides the round trip time parameter for window management. If the configured **round-trip-time** is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

The **min-available-bandwidth** parameter and the measured RTT together determine the threshold belo w which TCP aggressively maintains a window size sufficient to transmit at minimum available bandwi dth.

The **max-bandwidth-mbps** parameter and the measured RTT together determine the maximum windo w size.

Note Set the maximum bandwidth to match the worst-case bandwidth available on the physical link, keeping in mind other traffic that might be going across this link (for example, other FCIP tunnels, WAN limitati ons)–in other words, maximum bandwidth should be the total bandwidth minus all other traffic going ac ross that link.

To configure window management, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10`	Configures the maximum available bandwi dth at 900 Mbps, the minimum slow start threshold at 300 Mbps, and the RTT at 10 msec.
	`switch(config-profile)# no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10`	Reverts to the factory defaults. The FCIP defaults are maximum bandwidth at 1 Gbps , minimum available bandwidth at 500 Mbps, and RTT at 1 msec.
	`switch(config-profile)# tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200`	Configures the maximum available bandwi dth at 2000 Kbps, the minimum available bandwidth at 2000 Kbps, and the RTT at 200 msec.

Monitoring Congestion

By enabling the congestion window monitoring (CWM) parameter, you allow TCP to monitor congesti on after each idle period. The CWM parameter also determines the maximum burst size allowed after an idle period. By default, this parameter is enabled and the default burst size is 50 KB.

The interaction of bandwidth parameters and CWM and the resulting TCP behavior is outlined as follo ws:

- If the average rate of the Fibre Channel traffic over the preceding RTT is less than the min-avai lable-bandwidth multiplied by the RTT, the entire burst is sent immediately at the min-avai lable-bandwidth rate, provided no TCP drops occur.
- If the average rate of the Fibre Channel traffic is greater than min-available-bandwidth multiplied by the RTT, but less than max-bandwidth multiplied by the RTT, then if the Fibre Channel traffic is transmitted in burst sizes smaller than the configured CWM value the entire burst is sent immedi ately by FCIP at the max-bandwidth rate.
- If the average rate of the Fibre Channel traffic is larger than the min-available-bandwidth multipli ed by the RTT and the burst size is greater than the CWM value, then only a part of the burst is sent immediately. The remainder is sent with the next RTT.

The software uses standard TCP rules to increase the window beyond the one required to maintain the min-available-bandwidth to reach the max-bandwidth.

Note The default burst size is 50 KB.

Tip We recommend that this feature remain enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

To change the CWM defaults, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp cwm`	Disables congestion monitoring.
	`switch(config-profile)# tcp cwm`	Enables congestion monitoring and sets the burst size to its default size.
	`switch(config-profile)# tcp cwm burstsize 30`	Changes the burst size to 30 KB. The valid range is from 10 to 100 KB.
	`switch(config-profile)# no tcp cwm burstsize 25`	Leaves the CWM feature in an enabled state but changes the burst size to its factory default.

Estimating Maximum Jitter

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

You can configure the maximum estimated jitter in microseconds by the packet sender. The estimated variation should not include network queuing delay. By default, this parameter is enabled in Cisco MDS switches when IPS modules or MPS-14/2 modules are present.

The default value is 1000 microseconds for FCIP interfaces.

To configure the maximum jitter value, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# no tcp max-jitter`	Disables delay jitter estimation.
	`switch(config-profile)# tcp max-jitter`	Enables the delay jitter feature and sets the time to its factory default.
	`switch(config-profile)# tcp max-jitter 300`	Changes the time to 300 microseconds. The valid range is from 0 to 10000 microseconds.
	`switch(config-profile)# no tcp max-jitter 2500`	Leaves the delay jitter feature in an enabled state but changes the time to its factory default (1000 microseconds for FCIP interfaces).

Buffer Size

You can define the required additional buffering—beyond the normal send window size —that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default FCIP buffer size is 0 KB.

Note Use the default if the FCIP traffic is passing through a high throughput WAN link. If you have a mismatch in speed between the Fibre Channel link and the WAN link, then time stamp errors occur in the DMA bridge. In such a situation, you can avoid time stamp errors by increasing the buffer size.

To set the buffer size, follow these steps:

	Command	Purpose
Step 1	`switch(config-profile)# tcp send-buffer-size 5000`	Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 16384 KB.
	`switch(config-profile)# no tcp send-buffer-size 5000`	Reverts the switch to its factory default. The default is 0 KB.

Displaying FCIP Profile Configuration Information

Use the **show fcip profile** command to display FCIP profile configuration information.

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

Advanced FCIP Interface Configuration

This section describes the options you can configure on an FCIP interface to establish connection to a peer and includes the following topics:

To establish a peer connection, you must first create the FCIP interface and enter the `config-if` submode.

To enter the `config-if` submode, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal`	Enters configuration mode.
Step 2	`switch(config)# interface fcip 100`	Creates an FCIP interface (100).

Configuring Peers

To establish an FCIP link with the peer, you can use one of two options:

- Peer IP address—Configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- Special frames—Configures one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the switch WWN (sWWN) and profile ID along with the IP address.

Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can als o specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection. You can specify an IPv4 address or an IPv6 address.

To assign the peer information based on the IPv4 address and port number, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# peer-info ipaddr 10.1.1.1`	Assigns an IPv4 address to configure the peer i nformation. Because no port is specified, the default port number (3225) is used.
	`switch(config-if)# no peer-info ipaddr 10.10.1.1`	Deletes the assigned peer port information.

	Command	Purpose
Step 2	`switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000`	Assigns the IPv4 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.
	`switch(config-if)# no peer-info ipaddr 10.1.1.1 port 3000`	Deletes the assigned peer port information.
Step 3	`switch(config-if)# no shutdown`	Enables the interface.

To assign the peer information based on the IPv6 address and port number, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# peer-info ipaddr`	Assigns an IPv6 address to configure the peer i nformation. Because no port is specified, the default port number (3225) is used.
	`switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a`	Deletes the assigned peer port information.
Step 2	`switch(config-if)# peer-info ipaddr 2001:0db8:800:200c::417a port 3000`	Assigns the IPv6 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.
	`switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a port 3000`	Deletes the assigned peer port information.
Step 3	`switch(config-if)# ipv6 enable`	Enables IPv6 processing on the interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Special Frames

You can alternatively establish an FCIP link with a peer using an optional protocol called *special frames* . When special frames are enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exc hanged to discover and authenticate the link.

By default, the special frame feature is disabled. You must enable special frames on the interfaces on bo th peers to establish the FCIP link.

Note Refer to the Fibre Channel IP standards for further information on special frames.

Tip Special frame negotiation provides an additional authentication security mechanism because the link vali dates the WWN of the peer switch.

To enable special frames, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00`	Enables special frames and sets the peer WWN as specified. **Note** The peer WWN is the WWN of the peer switch. Use the **show wwn switch** co mmand to obtain the peer WWN.
	`switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00`	Disables special frames (default).
Step 2	`switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155`	Enables special frames and sets the peer WWN and the profile ID (155).
	`switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155`	Disables special frames (default).
Step 3	`switch(config-if)# no shutdown`	Enables the interface.

Active Connections

You can configure the required mode for initiating a TCP connection. By default, active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection rather waits for the peer to connect to it.

Note Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

To enable the passive mode, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# passive-mode`	Enables passive mode while attempting a TCP connection.
	`switch(config-if)# no passive-mode`	Reverts to the factory set default of using the active mode while attempting the TCP connection.
Step 2	`switch(config-if)# no shutdown`	Enables the interface.

Number of TCP Connections

You can specify the number of TCP connections from an FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure one or two TCP connections. For example, the Cisco PA-FC-1G Fibre Channel port adapter, which has only one (1) TCP connection, interoperates wi

th any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, then the software handles it and proceeds with just one connection.

To specify the TCP connection attempts, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# tcp-connection 1`	Specifies the number of TCP connections. Valid values are 1 or 2.
	`switch(config-if)# no tcp-connection 1`	Reverts to the factory set default of two attempts.
Step 2	`switch(config-if)# no shutdown`	Enables the interface.

Time Stamp Control

You can instruct the switch to discard packets that are outside the specified time. When enabled, this feature specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped.

By default, time stamp control is disabled in all switches in the Cisco MDS 9000 Family. If a packet arri ves within a 2000 millisecond interval (+ or –2000 msec) from the network time, that packet is accepted.

Note The default value for packet acceptance is 2000 microseconds.
If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the "NTP Confi guration" section on page 5-19).

Tip Do not enable time stamp control on an FCIP interface that has tape acceleration or write acceleration c onfigured.

To enable or disable the time stamp control, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# time-stamp` `Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link`	Enables time stamp checking for received packets with a default acceptable time difference of 2000 msec.
	`switch(config-if)# no time-stamp`	Disables (default) time stamps.
Step 2	`switch(config-if)# time-stamp acceptable-diff 4000`	Configures the packet acceptance time. The valid range is from 500 to 10,000 msec.
	`switch(config-if)# no time-stamp acceptable-diff 500`	Deletes the configured time difference and reverts the di fference to factory defaults. The default difference is a 2000-millisecond interval from the network time.
Step 3	`switch(config-if)# no shutdown`	Enables the interface.

B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. Figure 40-11 shows a typical SAN extension over an IP network.

Figure 40-11 FCIP B Port and Fibre Channel E Port

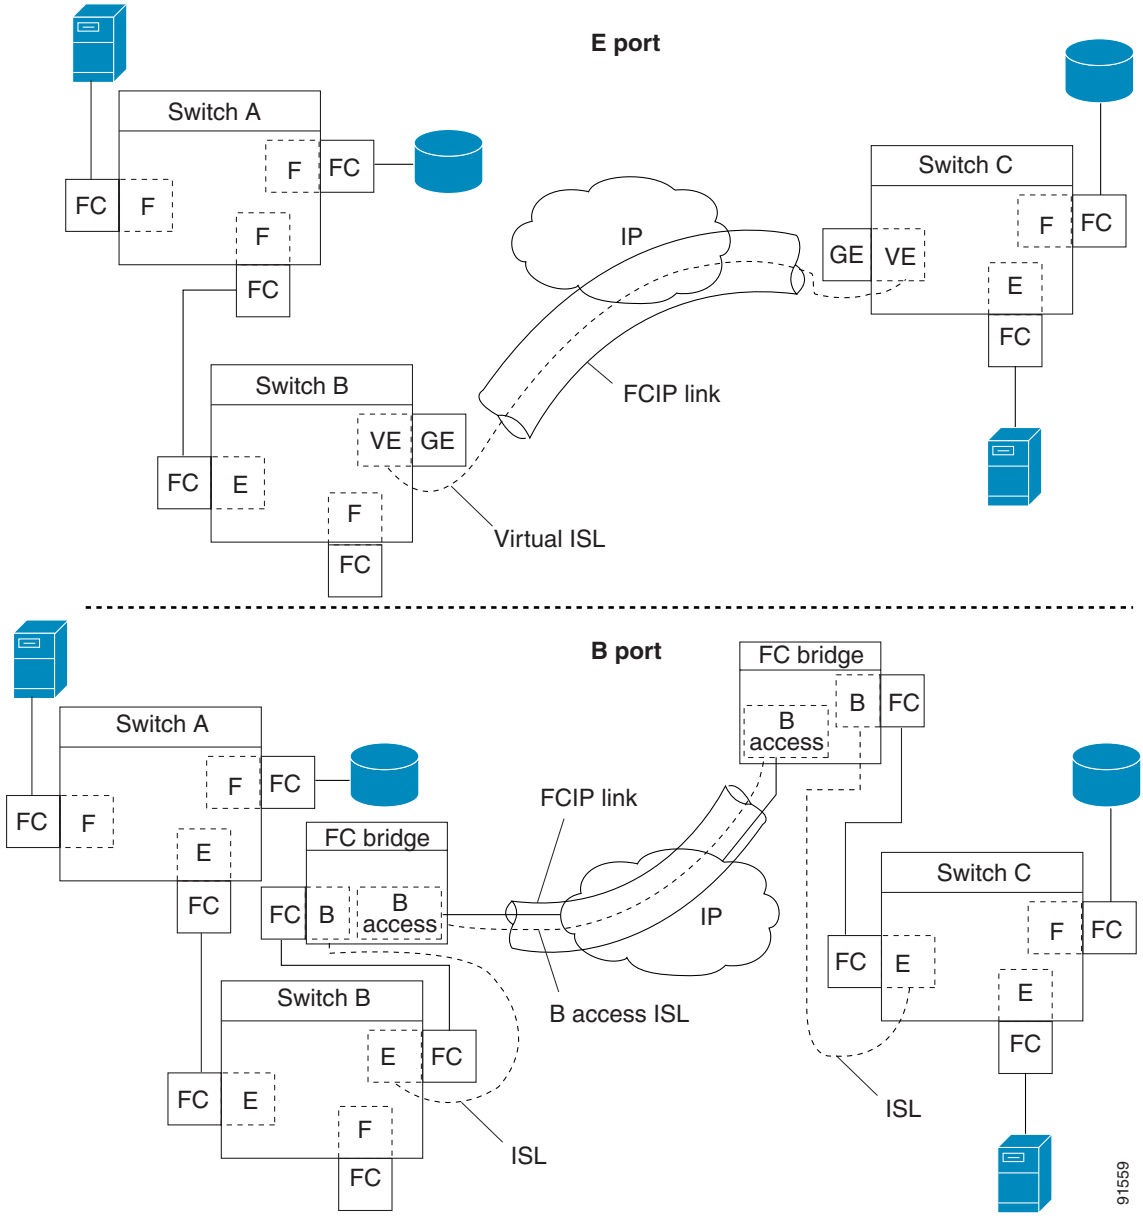

B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not interac

t with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over an FCIP link, B ports use a B access ISL.*

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corres ponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see Figure 40-12).

Figure 40-12 FCIP Link Terminating in a B Port Mode

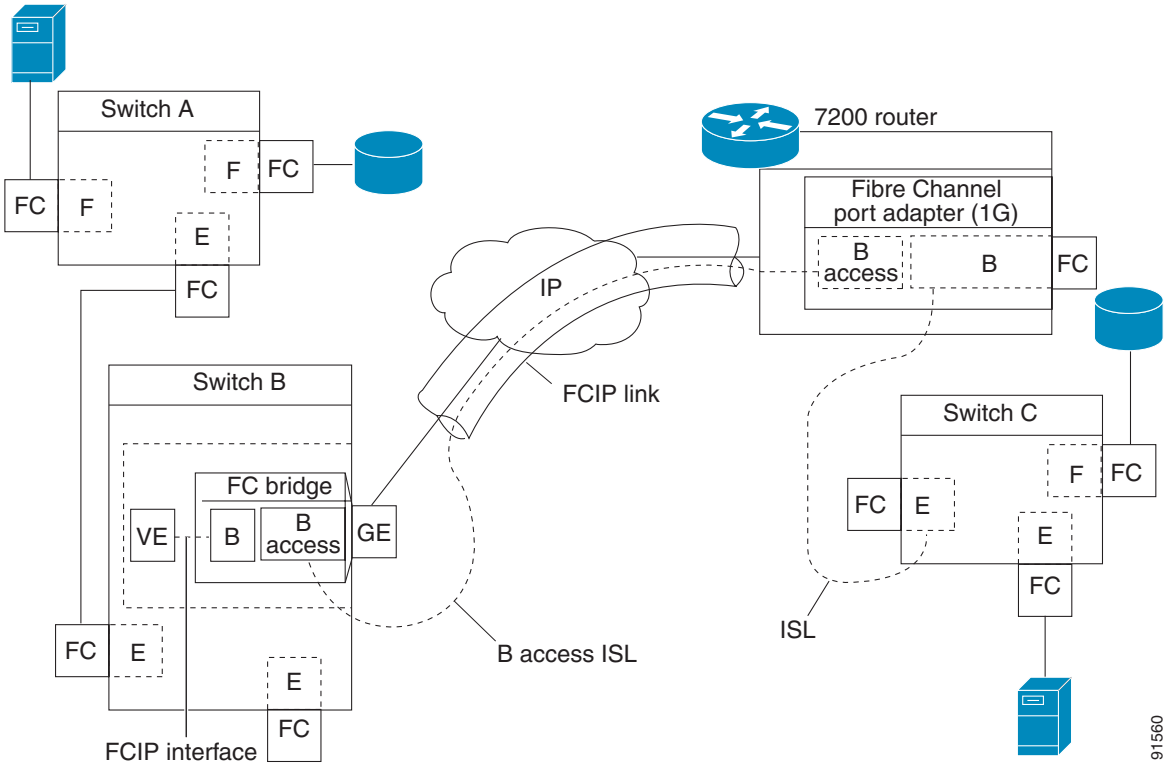

The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to co mmunicate directly with a Cisco MDS 9000 Family switch, eliminating the need for local bridge devices.

Configuring B Ports

When an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To enable B port mode, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# bport`	Enables B port mode on the FCIP interface.
	`switch(config-if)# no bport`	Reverts to E port mode on the FCIP interface (default).
Step 2	`switch(config-if)# bport-keepalive`	Enables the reception of keepalive responses sent by a remote peer.
	`switch(config-if)# no bport-keepalive`	Disables the reception of keepalive responses sent by a remote peer (default).

Quality of Service

The quality of service (QoS) parameter specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that co nnection.

To set the QoS values, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# qos control 24 data 26`	Configures the control TCP connection and data co nnection to mark all packets on that DSCP value. The control and data value ranges from 0 to 63.
	`switch(config-if)# no qos control 24 data 26`	Reverts the switch to its factory default (marks all c ontrol and data packets with DCSP value 0).

Configuring E Ports

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN (see Chapter 19, “Configuring and Managing VSANs”).
- Trunk mode and trunk allowed VSANs (see Chapter 15, “Configuring Trunking”).
- PortChannels (see Chapter 16, “Configuring PortChannels”):
  - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
  - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF (see Chapter 25, “Configuring Fibre Channel Routing Services and Protocols”).

- Fibre Channel domains (fcdomains) (see Chapter 17, "Configuring Domain Parameters.").
- Importing and exporting the zone database from the adjacent switch (see Chapter 23, "Configuring and Managing Zones").

Displaying FCIP Interface Information

Use the **show interface** commands to view the summary, counter, description, and status of the FCIP li nk. Use the output of these commands to verify the administration mode, the interface status, the operati onal mode, the related VSAN ID, and the profile used. See Example 40-3 through Example 40-6.

Example 40-3 Displays the FCIP Summary

```
switch# show fcip summary
--------------------------------------------------------------------------------
Tun prof    Eth-if    peer-ip       Status T W T Enc Comp  Bandwidth   rtt
                                           E A A            max/min    (us)
--------------------------------------------------------------------------------
10  91   GE4/1      3.3.3.2          UP   N N N  N    N  1000M/1000M  2000
11  11   GE3/1.601  30.1.1.2       DOWN   N N N  N    N   1000M/500M  1000
12  12   GE3/1.602  30.1.2.2       DOWN   N N N  N    N   1000M/500M  1000
13  0               0.0.0.0        DOWN   N N N  N    N
14  0               0.0.0.0        DOWN   N N N  N    N
15  0               0.0.0.0        DOWN   N N N  N    N
16  0               0.0.0.0        DOWN   N N N  N    N
17  0               0.0.0.0        DOWN   N N N  N    N
18  0               0.0.0.0        DOWN   N N N  N    N
19  0               0.0.0.0        DOWN   N N N  N    N
20  92   GE4/2      3.3.3.1          UP   N N N  N    N  1000M/1000M  2000
21  21   GE3/2.601  30.1.1.1       DOWN   N N N  N    N   1000M/500M  1000
22  22   GE3/2.602  30.1.2.1       DOWN   N N N  N    N   1000M/500M  1000
```

Example 40-4 Displays the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 10
fcip10 is up
    Hardware is GigabitEthernet
    Port WWN is 20:d0:00:0c:85:90:3e:80
    Peer port WWN is 20:d4:00:0c:85:90:3e:80
    Admin port mode is auto, trunk mode is on
    Port mode is E, FCID is 0x720000
    Port vsan is 91
    Speed is 1 Gbps
    Using Profile id 91  (interface GigabitEthernet4/1)
    Peer Information
      Peer Internet address is 3.3.3.2 and port is 3225
    Write acceleration mode is off
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is 256 KBytes
    IP Compression is disabled
    Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
    TCP Connection Information
      50529025 Active TCP connections
         Local 0.0.0.7:6, Remote 0.0.0.200:0
```

```
      0 host table full 0 target entries in use
      211419104 Attempts for active connections, 1500 close of connections
    TCP Parameters
      Path MTU 124160 bytes
      Current retransmission timeout is 124160 ms
      Round trip time: Smoothed 127829 ms, Variance: 14336
      Advertized window: Current: 0 KB, Maximum: 14 KB, Scale: 14336
      Peer receive window: Current: 0 KB, Maximum: 0 KB, Scale: 51200
      Congestion window: Current: 14 KB, Slow start threshold: 49344 KB
      Current Send Buffer Size: 206463 KB, Requested Send Buffer Size: 429496728
3 KB
      CWM Burst Size: 49344 KB
    5 minutes input rate 491913172779207224 bits/sec, 61489146597400903 bytes/se
c, 0 frames/sec
    5 minutes output rate 491913175298921320 bits/sec, 61489146912365165 bytes/s
ec, 14316551 frames/sec
      5702 frames input, 482288 bytes
         5697 Class F frames input, 481736 bytes
         5 Class 2/3 frames input, 552 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      5704 frames output, 482868 bytes
         5698 Class F frames output, 482216 bytes
         6 Class 2/3 frames output, 652 bytes
         0 Error frames
```

Example 40-5 Displays Detailed FCIP Interface Standard Counter Information

```
switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    5 minutes input rate 207518944 bits/sec, 25939868 bytes/sec, 12471 frames/sec
    5 minutes output rate 205340328 bits/sec, 25667541 bytes/sec, 12340 frames/sec
      2239902537 frames input, 4658960377152 bytes
         18484 Class F frames input, 1558712 bytes
         2239884053 Class 2/3 frames input, 4658958818440 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      2215051484 frames output, 4607270186816 bytes
         18484 Class F frames output, 1558616 bytes
         2215033000 Class 2/3 frames output, 4607268628200 bytes
         0 Error frames
```

Example 40-6 Displays the FCIP Interface Description

```
switch# show interface fcip 51 description
FCIP51
    Sample FCIP interface
```

The txbytes is the amount of data before compression. After compression, the compressed txbytes bytes are transmitted with compression and the uncompressed txbytes bytes are transmitted without compres sion. A packet may be transmitted without compression, if it becomes bigger after compression (see Example 40-7).

Example 40-7 Displays Brief FCIP Interface Counter Information

```
switch# show interface fcip 3 counters brief
--------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
```

```
                    ----------------------------  ----------------------------
                    Rate     Total                Rate     Total
                    Mbits/s  Frames               Mbits/s  Frames
-------------------------------------------------------------------------------
fcip3               9        0                    9        0
```

Advanced FCIP Features

You can significantly improve application performance by configuring one or more of the following opti ons for the FCIP interface.

FCIP Write Acceleration

The FCIP write acceleration feature enables you to significantly improve application write performanc e when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.

Note The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel the write acceleration feature will be turned operati onally off.

In Figure 40-13, the WRITE command without write acceleration requires two round trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latenc y over the FCIP link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Figure 40-13 FCIP Link Write Acceleration

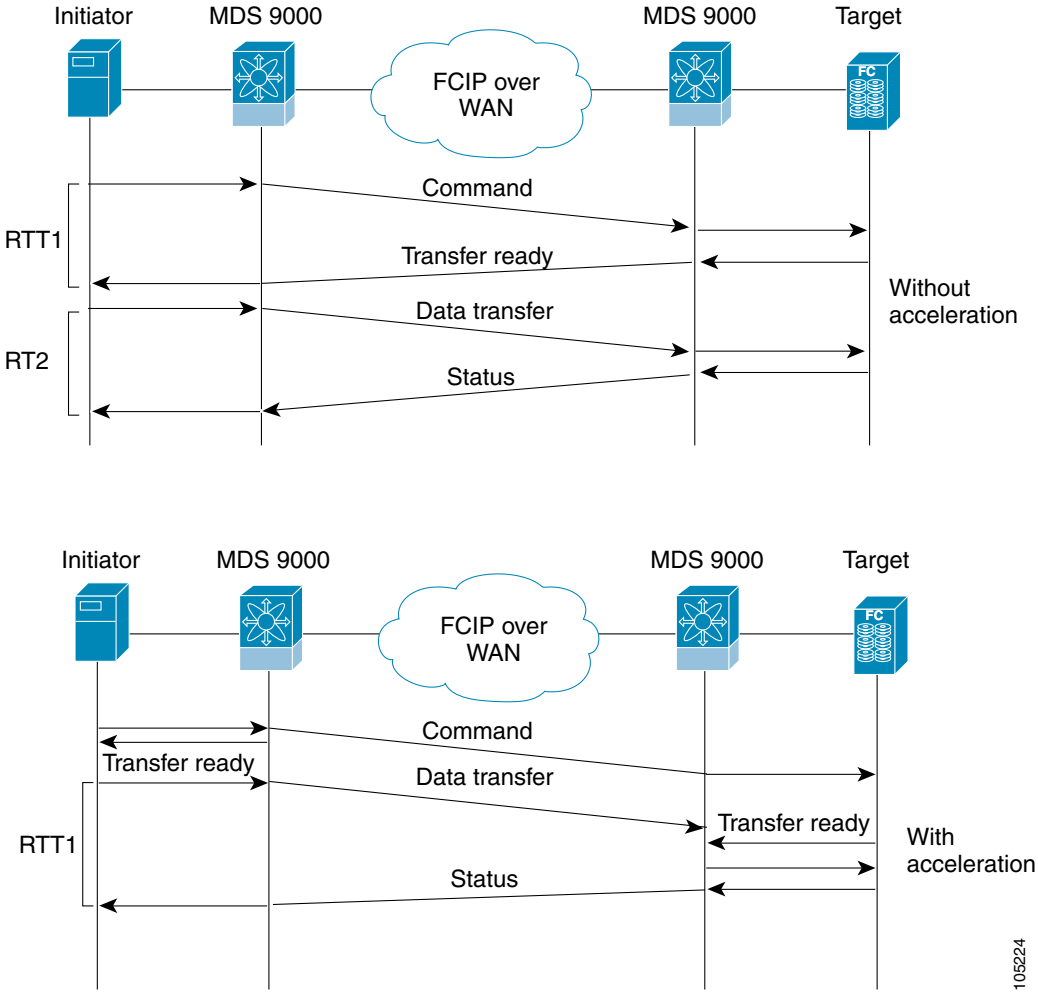

Tip FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel configured with channel mode active. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or failed WRITE or READ operations.

Tip Do not enable time stamp control on an FCIP interface with write acceleration configured.

Note Write acceleration cannot be used across FSPF equal cost paths in FCIP deployments. Native Fibre Channel write acceleration can be used with Port Channels. Also, FCIP write acceleration can be used in Port Channels configured with channel mode active or constructed with Port Channel Protocol (PCP.

Caution FCIP write acceleration with FCIP ports as members of PortChannels in Cisco MDS SAN-OS Release 2.0(1b) and later are incompatible with the FCIP write acceleration in earlier releases.

Configuring FCIP Write Acceleration

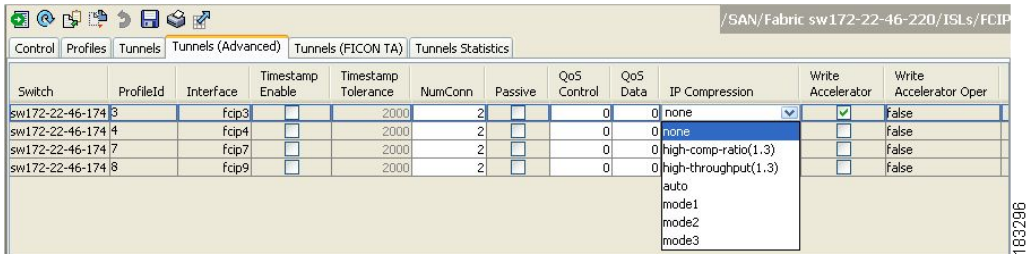

To enable write acceleration, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 51` `switch1(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch1(config-if)# write-accelerator`	Enables write acceleration.
	`switch1(config-if)# no write-accelerator`	Disables write acceleration (default).

Displaying Write Acceleration Activity Information

Example 40-8 through Example 40-10 show how to display information about write acceleration activi ty.

Example 40-8 Displays Exchanges Processed by Write Acceleration at the Specified Host End FCIP Li nk.

```
switch# show fcip host-map 100

 MAP TABLE (5 entries  TOTAL entries 5)

 OXID | RXID | HOST FCID| TARG FCID| VSAN |  Index
------+------+----------+----------+------+---------
0xd490|0xffff|0x00690400|0x00620426|0x0005|0x0000321f
0xd4a8|0xffff|0x00690400|0x00620426|0x0005|0x00003220
0xd4c0|0xffff|0x00690400|0x00620426|0x0005|0x00003221
0xd4d8|0xffff|0x00690400|0x00620426|0x0005|0x00003222
0xd4f0|0xffff|0x00690400|0x00620426|0x0005|0x00003223
```

Example 40-9 Displays Exchanges Processed by Write Acceleration at the Specified Target End FCIP Li nk

```
switch# show fcip target-map 100

 MAP TABLE (3 entries TOTAL entries 3)
```

```
 OXID | RXID | HOST FCID| TARG FCID| VSAN |  Index
------+------+----------+----------+------+---------
0xc308|0xffff|0x00690400|0x00620426|0x0005|0x00003364
0xc320|0xffff|0x00690400|0x00620426|0x0005|0x00003365
0xc338|0xffff|0x00690400|0x00620426|0x0005|0x00003366
```

Example 40-10 Displays Detailed FCIP Interface Write Acceleration Counter Information, if Enabled

```
switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    Write Accelerator statistics
      6091 packets in      5994 packets out
      0 frames dropped  0 CRC errors
      0 rejected due to table full
      0 ABTS sent       0 ABTS received
      0 tunnel synchronization errors
      37 writes recd     37 XFER_RDY sent (host)
      0 XFER_RDY rcvd (target)
      37 XFER_RDY rcvd (host)
      0 XFER_RDY not proxied due to flow control (host)
      0 bytes queued for sending
      0 estimated bytes queued on the other side for sending
      0 times TCP flow ctrl(target)
      0 bytes current TCP flow ctrl(target)
```

FCIP Tape Acceleration

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS SAN-OS provides both tape write and read acceleration.

Applications that access tape drives normally have only one SCSI WRITE or READ operation outstandi ng to it. This single command process limits the benefit of the tape acceleration feature when using an FCIP tunnel over a long-distance WAN link. It impacts backup, restore, and restore performance becaus e each SCSI WRITE or READ operation does not complete until the host receives a good status respons e from the tape drive. The FCIP tape acceleration feature helps solve this problem. It improves tape bac kup, archive, and restore operations by allowing faster data streaming between the host and tape drive o ver the WAN link.

In an example of tape acceleration for write operations, the backup server in Figure 40-14 issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This res ponse allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data witho ut proxying. The proxy method improves the performance on WAN links.

Figure 40-14 FCIP Link Tape Acceleration for Write Operations

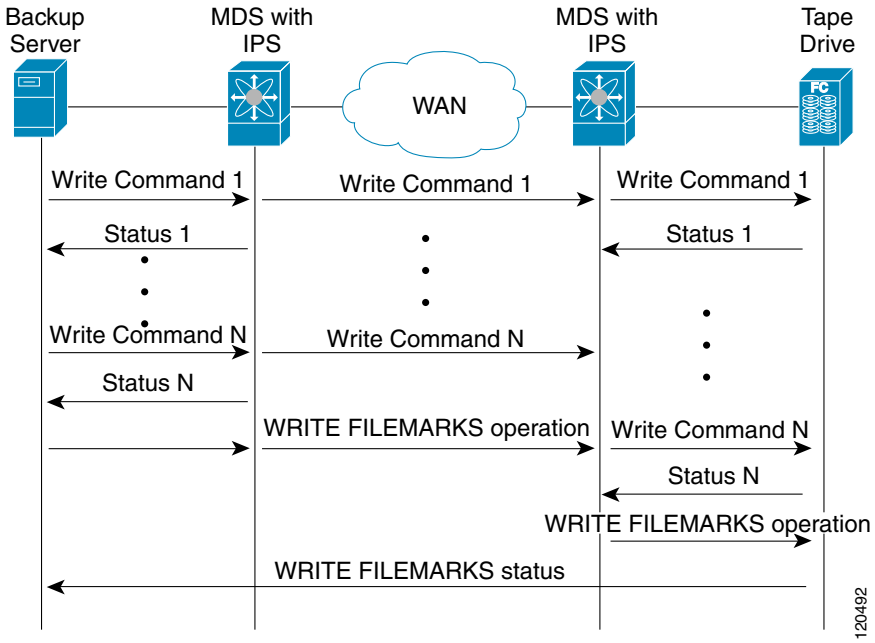

At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has rec eived. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape dri ve before forwarding the data.

Note In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape li brary environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these ses sions. The workaround is to keep the FCIP link disabled for a couple of minutes before enabli ng the link. Note that this does not apply to tape environments where the tape drives are either direct FC attached or exported as LUN 0.

The Cisco SAN-OS provides reliable data delivery to the remote tape drives using TCP/IP over the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data wi th the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco SAN-OS software.

In an example of tape acceleration for read operations, the restore server in Figure 40-15 issues read o perations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI read operations from the host, sends out SCSI read operations o n its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI read operations from the host, sends out the cached data. This method results in more data being sent over the FCIP tunnel in the same time period compared to the ti me taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Figure 40-15 FCIP Link Tape Acceleration for Read Operations

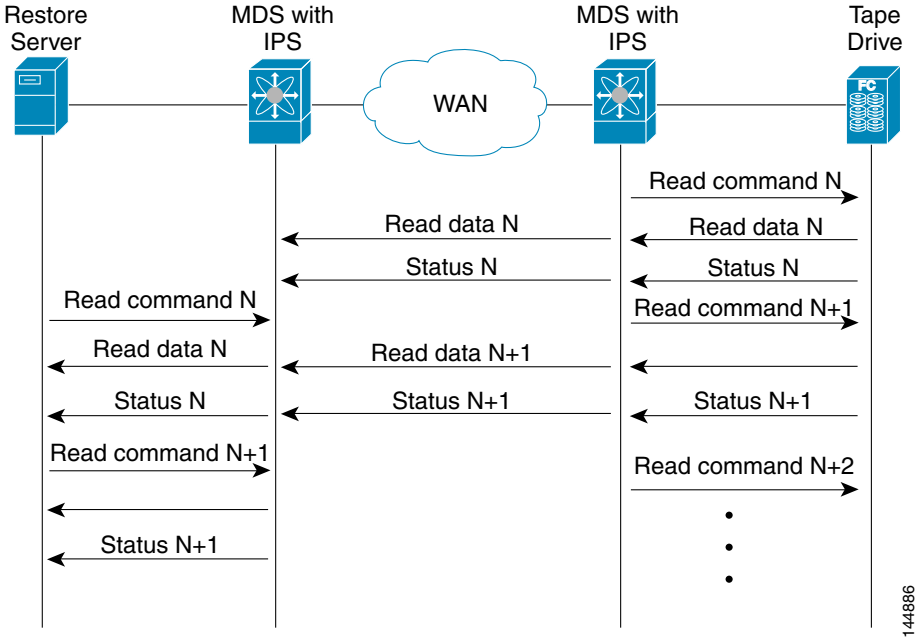

The Cisco SAN-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco SAN-OS software recovers from any other errors.

Note The FCIP tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tape acceleration feature is turned operatio nally off.

Tip FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.

Caution When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that i nterface. Likewise, if an FCIP interface is up in a FICON VSAN, tape acceleration cannot be enabled o n that interface.

Note When you enable the tape acceleration feature for an FCIP tunnel, the tunnel is reinitialized and the wri te and read acceleration feature is also automatically enabled.

In tape acceleration for writes, after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch by not pro xying the Transfer Ready. On completion of a write operation when some data buffers are freed, the loc al Cisco MDS switch resumes the proxying. Likewise, in tape acceleration for reads, after a certain amo

unt of data has been buffered at the local Cisco MDS switch, the read operations to the tape drive are flo w controlled by the remote Cisco MDS switch by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow contro l buffer size (the maximum buffer size is 12 MB).

Tip We recommend that you use the default option for flow- control buffering.

Tip Do not enable time- stamp control on an FCIP interface with tape acceleration configured.

Note If one end of the FCIP tunnel is running Cisco MDS SAN-OS Release 3.0(1) or later, and the other end is running Cisco MDS SAN-OS Release 2.x, and tape acceleration is enabled, then the FCIP tunnel wi ll run only tape write acceleration, not tape-read acceleration.

Tape Library LUN Mapping for FCIP Tape Acceleration

If a tape library provides logical unit (LU) mapping and FCIP tape acceleration is enabled, you must as sign a unique LU number (LUN) to each physical tape drive accessible through a target port.

Figure 40-16 shows tape drives connected to Switch 2 through a single target port. If the tape library pro vides LUN mapping, then all the four tape drives should be assign unique LUNs.

Figure 40-16 FCIP LUN Mapping Example

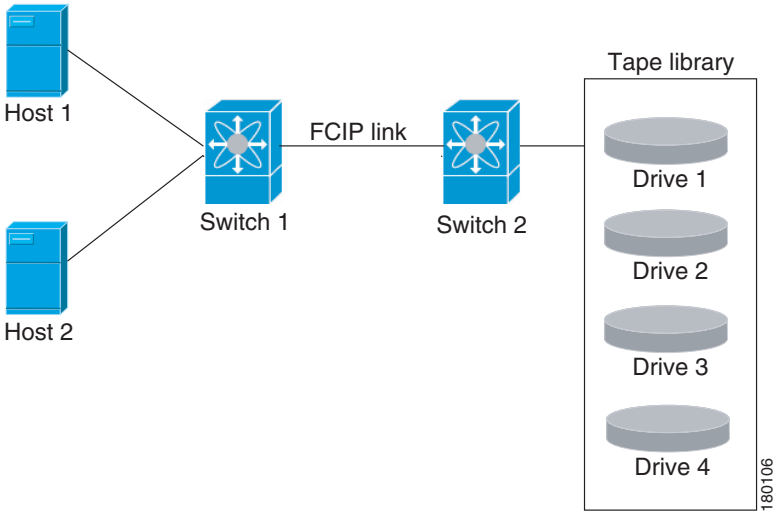

For the mappings described in Table 40-1 and Table 40-2, Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

Table 40-1 describes correct tape library LUN mapping.

Table 40-1 Correct LUN Mapping Example with Single Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 3	Drive 3
	LUN 4	Drive 4

Table 40-2 describes incorrect tape library LUN mapping.

Table 40-2 Incorrect LUN Mapping Example with Single Hosts Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 1	Drive 3
	LUN 2	Drive 4

Another example setup is when a tape drive is shared by multiple hosts through a single tape port. For i nstance, Host 1 has access to Drive1 and Drive2, and Host 2 has access to Drive 2, Drive 3, and Drive 4. A correct LUN mapping configuration for such a setup is shown in Table 40-3.

Table 40-3 Correct LUN Mapping Example with Multiple Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 2	Drive 2
	LUN 3	Drive 3
	LUN 4	Drive 4

Configuring FCIP Tape Acceleration

To enable FCIP tape acceleration, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch1(config)# interface fcip 5` `switch1(config-if)#`	Creates an FCIP interface (5).

	Command	Purpose
Step 3	`switch1(config-if)# write-accelerator tape-accelerator`	Enables tape acceleration (and write acceleration—if not already enabled).
	`switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size auto`	Enables tape acceleration with automatic flow control (default).
	`switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size 2048`	Sets tape acceleration flow control buffer size to 2 MB.
	`switch1(config-if)# no write-accelerator tape-accelerator`	Disables tape acceleration (default) and resets the FCIP tunnel. **Note** The write acceleration feature remains enabled.
	`switch1(config-if)# no write-accelerator tape-accelerator flow-control-buffer-size 2048`	Changes the flow control buffer size to the default value of automatic. The tape acceleration and write acceleration features remain enabled. This command does not reset the FCIP tunnel.
	`switch1(config-if)# no write-accelerator`	Disables both the write acceleration and tape acceleration features and resets the FCIP tunnel.

Displaying Tape Acceleration Activity Information

Example 40-11 through Example 40-14 show how to display information about tape acceleration activity.

Example 40-11 Displays Information About Tapes for Which Exchanges are Tape Accelerated

```
switch# show fcip tape-session summary
--------------------------------------------------------------------------------
Tunnel    Tunnel End    tape-fcid    lun        vsan    num-hosts
--------------------------------------------------------------------------------
  1       host-end       EF0001      0x0002     0001      1
  2       targ-end       650001      0x0003     0010      2
--------------------------------------------------------------------------------
```

Example 40-12 Displays Information About Tapes for Which Exchanges are Tape Accelerated at the Host-End FCIP Link

```
switch# show fcip tape-session tunnel 1 host-end

HOST TAPE SESSIONS (1 entries TOTAL entries 1)

Host Tape Session #1
 FCID 0xEF0001, VSAN 1, LUN 0x0002
 Outstanding Exchanges 0, Outstanding Writes 0
 Target End Write Buffering 0 Bytes, Auto Max Writes 3
 Flags 0x0, FSM state Non TA Mode
 Cached Reads 0
 First index 0xfffffff7, Last index 0xfffffff7, RA index 0x0000f99a
 Current index=0xfffffffe, Els Oxid 0xfff7
 Hosts 1
  FCID 0x770100
```

Example 40-13 Displays Information About Tapes for Which Exchanges are Tape Accelerated at the Target-End FCIP Link

```
switch# show fcip tape-session tunnel 1 targ-end

TARGET TAPE SESSIONS (1 entries TOTAL entries 1)

Target Tape Session #1
  FCID 0xEF0001, VSAN 1, LUN 0x0002
  Outstanding Exchanges 0, Outstanding Writes 0
  Host End Read Buffering 0 Bytes, Auto Max Read Blocks 3
  Flags 0x800, Timer Flags 0x0
  FSM State Default, Prev FSM State Bypass
  Relative Block offset 0
  First index 0xfffffff7, Last index 0xfffffff7, RA index 0x0000f99a
  Current index=0xfffffffe, Els Oxid 0xfff7
  Hosts 1
   FCID 0x770100
```

Example 40-14 Displays Detailed FCIP Interface Tape Acceleration Counter Information, if Enabled

```
switch# show interface fcip 1 counters
fcip1
   TCP Connection Information
....
   Tape Accelerator statistics
      1 Host Tape Sessions
      0 Target Tape Sessions
     Host End statistics
       Received 31521 writes, 31521 good status, 0 bad status
       Sent 31517 proxy status, 4 not proxied
       Estimated Write buffer 0 writes 0 bytes
       Received 31526 reads, 10 status
       Sent 31516 cached reads
       Read buffer 0 reads, 0 bytes
     Host End error recovery statistics
       Sent REC 0, received 0 ACCs, 0 Rejects
       Sent ABTS 0, received 0 ACCs
       Received 31 RECs, sent 2 ACCs, 0 Rejects
       Received 0 SRRs, sent 0 ACCs, 0 Rejects
       Received 0 TMF commands
     Target End statistics
       Received 0 writes, 0 good status, 0 bad status
       Write Buffer 0 writes, 0 bytes
       Received 0 reads, 0 good status, 0 bad status
       Sent 0 reads, received 0 good status, 0 bad status
       Sent 0 rewinds, received 0 good status, 0 bad status
       Estimated Read buffer 0 reads, 0 bytes
     Target End error recovery statistics
       Sent REC 0, received 0 ACCs, 0 Rejects
       Sent SRR 0, received 0 ACCs
       Sent ABTS 0, received 0 ACCs
       Received 0 TMF commands
```

FCIP Compression

The FCIP compression feature allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled. When enabled, the software defaults to using the **auto** mode (if a mode is not specified).

You can configure FCIP compression using one of the following modes:

- **mode1** is a fast compression mode for high bandwidth links (> 25 Mbps).
- **mode2** is a moderate compression mode for moderately low bandwidth links (between 10 and 25 Mbps).
- **mode3** is a high compression mode for low bandwidth links (< 10 Mbps).
- **auto** (default) mode picks the appropriate compression scheme based on the bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters).

The IP compression feature behavior differs between the IPS module and the MPS-14/2 module—while **mode2** and **mode3** perform software compression in both modules, **mode1** performs hardware-based c ompression in MPS-14/2 modules and software compression in IPS modules.

Note The Cisco MDS 9216i Switch also supports the IP compression feature. The integrated supervisor mo dule has the same hardware components that are available in the MPS-14/2 module.

Caution The compression modes in Cisco SAN-OS Release 2.0(1b) and later are incompatible with the compres sion modes in Cisco SAN-OS Release 1.3(1) and earlier.

Tip While upgrading from Cisco SAN-OS Release 1.x to Cisco SAN-OS Release 2.0(1b) or later, we reco mmend that you disable compression before the upgrade procedure, and then enable the required mode after the upgrade procedure.

If both ends of the FCIP link are running Cisco SAN-OS Release 2.0(1b) or later and you enable compres sion at one end of the FCIP tunnel, be sure to enable it at the other end of the link.

Configuring FCIP Compression

To enable FCIP compression, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fcip 51` `switch(config-if)#`	Creates an FCIP interface (51).
Step 3	`switch(config-if)# ip-compression mode3`	Enables high compression for low bandwi dth links.
	`switch(config-if)# ip-compression mode3`	Defaults to using the **auto** mode.
	`switch(config-if)# no ip-compression`	Disables (default) the FCIP compression feature.

Displaying FCIP Compression Information

Example 40-15 and Example 40-16 show how to display FCIP compression information.

Example 40-15 Displays Detailed FCIP Interface Compression Information, if Enabled

```
switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    IP compression statistics
      208752 rxbytes, 208752 rxbytes compressed
      5143584 txbytes
         0 txbytes compressed, 5143584 txbytes non-compressed
         1.00 tx compression ratio
```

Example 40-16 Displays the Compression Engine Statistics for the MPS-14/2 Module

```
switch#  show ips stats hw-comp all
HW Compression Statistics for port GigabitEthernet3/1
    Compression stats
      0 input bytes,  0 output compressed bytes
      0 input pkts,    0 output compressed pkts
    Decompression stats
      0 input compressed bytes, 0 output bytes
      0 input compressed pkts,   0 output pkts
    Passthru stats
      0 input bytes, 0 output bytes
      0 input pkts,   0 output pkts
    Miscellaneous stats
      32 min input pktlen,  32 max input pktlen
      28 min output pktlen, 28 max output pktlen
      0 len mismatch,      0 incomplete processing
      0 invalid result,    0 invalid session drop
      0 comp expanded
HW Compression Statistics for port GigabitEthernet3/2
    Compression stats
      0 input bytes,  0 output compressed bytes
      0 input pkts,    0 output compressed pkts
    Decompression stats
      0 input compressed bytes, 0 output bytes
      0 input compressed pkts,   0 output pkts
    Passthru stats
      0 input bytes, 0 output bytes
      0 input pkts,   0 output pkts
    Miscellaneous stats
      32 min input pktlen,  32 max input pktlen
      28 min output pktlen, 28 max output pktlen
      0 len mismatch,      0 incomplete processing
      0 invalid result,    0 invalid session drop
      0 comp expanded
```

Default Settings

Table 40-4 lists the default settings for FCIP parameters.

Table 40-4 Default FCIP Parameters

Parameters	Default
TCP default port for FCIP	3225
minimum-retransmit-time	200 msec
Keepalive timeout	60 sec
Maximum retransmissions	4 retransmissions
PMTU discovery	Enabled
pmtu-enable reset-timeout	3600 sec
SACK	Enabled
max-bandwidth	1Gbps
min-available-bandwidth	500 Mbps
round-trip-time	1 msec
Buffer size	0 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
TCP connection mode	Active mode is enabled
special-frame	Disabled
FCIP timestamp	Disabled
acceptable-diff range to accept packets	+/– 2000 msec
B port keepalive responses	Disabled
Write acceleration	Disabled
Tape acceleration	Disabled

CHAPTER 41

Configuring the SAN Extension Tuner

The SAN extension tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. You can specify the size of the test I/O transfers and how many concurrent or serial I/Os to generate while testing. The SET reports the resulting I/Os per second (IOPS) and I/O latency, which helps you determine the number of concurrent I/Os needed to maximize FCIP throughput.

This chapter includes the following sections:

About the SAN Extension Tuner

Note SAN Extension Tuner is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. To achieve maximum throughput performance across the fabric, you can tune the following configuration parameters:

- The TCP parameters for the FCIP profile (see the "Window Management" section on page 40-14).
- The number of concurrent SCSI I/Os generated by the application.
- The transfer size used by the application over an FCIP link.

SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options (see Figure 41-1).

Figure 41-1 SCSI Command Generation to the Virtual Target

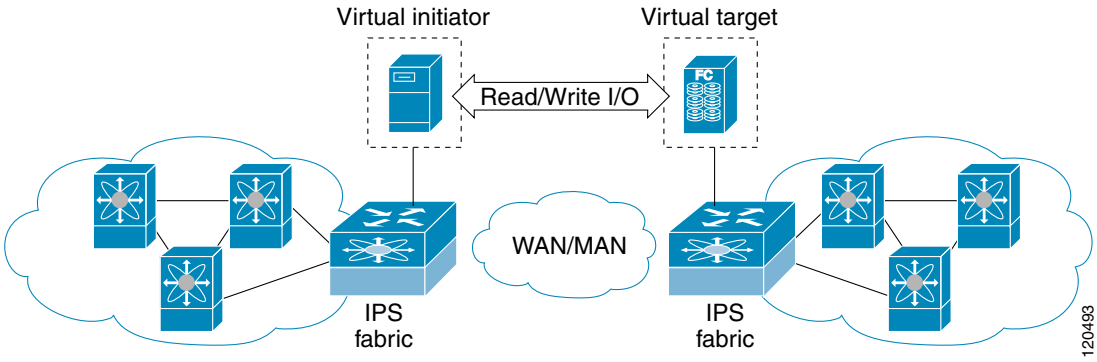

The SET feature assists with tuning by generating varying SCSI traffic workloads. It also measures throughput and response time per I/ O over an FCIP link.

Before tuning the SAN fabric, be aware of the following guidelines:

- Following these implementation details:
  - The tuned configuration is not persistent.
  - The virtual N ports created do not register FC4 features supported with the name server. This is to avoid the hosts in the SAN from discovering these N ports as regular initiators or targets.
  - Login requests from other initiators in the SAN are rejected.
  - The virtual N ports do not implement the entire SCSI suite; it only implements the SCSI read and write commands.
  - Tuner initiators can only communicate with tuner targets.
- Verify that the Gigabit Ethernet interface is up at the physical layer (GBIC and Cable connected—an IP address is not required).
- Enable iSCSI on the switch (no other iSCSI configuration is required).
- Create an iSCSI interface on the Gigabit Ethernet interface and enable the interface (no other iSCSI interface configuration is required)(see the "Creating iSCSI Interfaces" section on page 42-5).
- Configure the virtual N ports in a separate VSAN or zone as required by your network.
- Be aware that a separate VSAN with only virtual N ports is not required, but is recommended as some legacy HBAs may fail if logins to targets are rejected.
- Do not use same Gigabit Ethernet interface to configure virtual N ports and FCIP links—use different Gigabit Ethernet interfaces. While this is not a requirement, it is recommended as the traffic generated by the virtual N ports may interfere with the performance of the FCIP link.

SAN Extension Tuner Setup

Figure 41-2 provides a sample physical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 41-2 N Port Tuning Configuration Physical Example

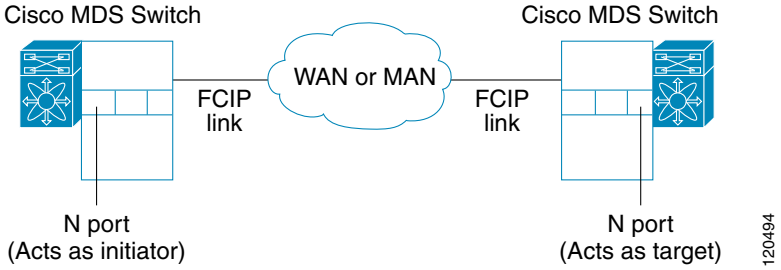

Figure 41-3 provides a sample logical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 41-3 Logical Example of N Port Tuning for a FCIP Link

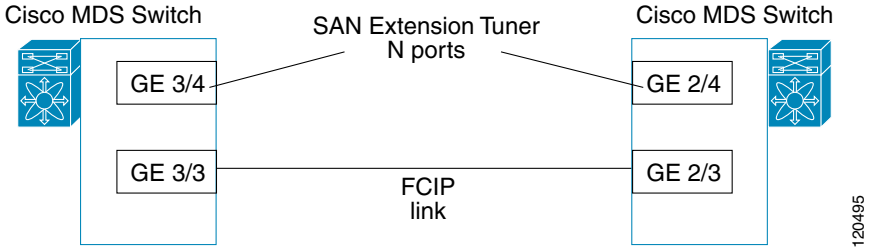

Data Pattern

By default, an all-zero pattern is used as the pattern for data generated by the virtual N ports. You can optionally specify a file as the data pattern to be generated by selecting a data pattern file from one of three locations: the bootflash: directory, the volatile: directory, or the slot0: directory. This option is especially useful when testing compression over FCIP links. You can also use Canterbury corpus or artificial corpus files for benchmarking purposes.

License Prerequisites

To use the SET, you need to obtain the SAN_EXTN_OVER_IP license (see Chapter 3, "Obtaining and Installing Licenses").

Configuring the SAN Extension Tuner

This section includes the following topics:

Tuning Guidelines

To tune the required FCIP link, follow these steps:

Step 1 Configure the nWWN for the virtual N ports on the switch.

Step 2 Enable iSCSI on the interfaces on which you want to create the N ports.

Step 3 Configure the virtual N ports on either side of the FCIP link.

Step 4 Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see Chapter 23, "Configuring and Managing Zones") or VSANs (see Chapter 19, "Configuring and Managing VSANs") to segregate the real initiators. Ensure that the zoning configuration is setup to allow the virtual N-ports to communicate with each other.

Step 5 Start the SCSI read and write I/Os.

Step 6 Add more N ports (as required) to other Gigabit Ethernet ports in the switch to obtain maximum throughput. One scenario that may require additional N ports is if you use FCIP PortChannels.

Tuner Initialization

The tuning feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, tuning is globally enabled for the entire switch.

To enable the tuning feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# san-ext-tuner enable`	Enables tuning.
	`switch(config)# no san-ext-tuner enable`	Removes the currently applied tuning configuration and disables tuning (default).

nWWN Configuration

To configure the nWWNs for the tuner in this switch, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.

Virtual N Port Configuration

To configure the virtual N port for tuning, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi enable`	Enables iSCSI.
Step 3	`switch(config)# interface iscsi 3/4` `switch(config-if)#`	Creates an iSCSI interface and enters interface configuration submode.
Step 4	`switch(config-if)# no shutdown`	Enables the iSCSI interface.
Step 5	`switch(config-if)# end` `switch#`	Returns to EXEC mode.
Step 6	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 7	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 8	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
	`switch(san-ext)# no nport pWWN 22:34:56:78:90:12:34:56 vsan 200 interface gigabitethernet 3/4`	Removes a virtual N port on the specified Gigabit Ethernet port and VSAN.

SCSI Read/Write Assignment

You can assign SCSI read and write commands on a one-time basis or on a continuous basis.

To assign SCSI read or write commands on a one-time basis, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 3	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	`switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000`	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the **read** command. The total number of I/Os is 5,000,000 bytes.
Step 5	`switch(san-ext-nport)# write command-id 101 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000`	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the **write** command received by the target. The total number of I/Os is 5,000,000 bytes.

	Command	Purpose
Step 6	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.
	`switch(san-ext-nport)# stop all`	Stops all outstanding commands.
Step 7	`switch(san-ext-nport)# clear counters`	Clears the counters associated with this N port.
Step 8	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

To generate SCSI read or write commands continuously, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 3	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	`switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous`	Configures SCSI commands to be read continuously. **Tip** Use the **stop command-id** command to stop the outstanding configuration.
Step 5	`switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous`	Configures SCSI commands to be written continuously.
Step 6	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.
	`switch(san-ext-nport)# stop command-id all`	Stops all outstanding commands.
Step 7	`switch(san-ext-nport)# clear counters`	Clears the counters associated with this N port.
Step 8	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

To specify a transfer ready size for a SCSI write command, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 3	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	`switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000`	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the **write** command received by the target. The total number of I/Os is 5,000,000 bytes.

	Command	Purpose
Step 5	`switch(san-ext-nport)# transfer-ready-size 512000`	Specifies the maximum transfer ready size of 512,000 bytes as a target for SCSI write commands. For a SCSI **write** command with a larger size, the target performs multiple transfers based on the specified transfer size.
	`switch(san-ext-nport)# no transfer-ready-size 512000`	Removes the specified transfer ready size configuration for SCSI write commands.
Step 6	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.
Step 7	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

SCSI Tape Read/Write Assignment

You can assign SCSI tape read and write commands on a one-time basis or on a continuous basis.

Note There is only one outstanding I/O at a time to the virtual N-port that emulates the tape behavior.

To assign SCSI tape read and or write commands on a one-time basis, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 3	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	`switch(san-ext-nport)# tape-read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 num-transactions 5000000 filemark-frequency 32`	Specifies a transfer size of 512,000 bytes with space over the filemark every 32 SCSI read commands. The total number of I/Os is 5,000,000 bytes.
Step 5	`switch(san-ext-nport)# tape-write command-id 101 target 22:22:22:22:22:22:22:22 transfer-size 512000 num-transactions 5000000 filemark-frequency 32`	Specifies a transfer size of 512,000 bytes with filemarks written every 32 SCSI write commands. The total number of I/Os is 5,000,000 bytes.
Step 6	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.
	`switch(san-ext-nport)# stop all`	Stops all outstanding commands.
Step 7	`switch(san-ext-nport)# clear counters`	Clears the counters associated with this N port.
Step 8	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

To generate SCSI tape read or write commands continuously, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nWWN 10:00:00:00:00:00:00:00`	Configures the nWWN for the SAN extension tuner.
Step 3	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	`switch(san-ext-nport)# tape-read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 continuous filemark-frequency 32`	Configures SCSI tape read commands to be issued continuously. **Tip** Use the **stop command-id** command to stop the outstanding configuration.
Step 5	`switch(san-ext-nport)# tape-write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 continuous filemark-frequency 32`	Configures SCSI tape write commands to be issued continuously.
Step 6	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.
	`switch(san-ext-nport)# stop command-id all`	Stops all outstanding commands.
Step 7	`switch(san-ext-nport)# clear counters`	Clears the counters associated with this N port.
Step 8	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

Configuring a Data Pattern

To optionally configure a data pattern for SCSI commands, follow these steps:

	Command	Purpose
Step 1	`switch# san-ext-tuner` `switch(san-ext)#`	Enters the SET configuration submode.
Step 2	`switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4` `switch(san-ext-nport)#`	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 3	`switch(san-ext-nport)# data-pattern-file bootflash://DataPatternFile`	Specifies the data pattern used by the N port to generate data as a target for **read** commands and initiator for **write** commands.
	`switch(san-ext-nport)# no data-pattern-file`	Removes the specified transfer ready size configuration for SCSI write commands and defaults to using the all-zero pattern.
Step 4	`switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000`	Specifies a transfer size of 512,000 bytes with two outstanding I/Os. The total number of I/Os is 5,000,000 bytes.
Step 5	`switch(san-ext-nport)# stop command-id 100`	Stops the command with the specified ID.

	Command	Purpose
Step 6	`switch(san-ext-nport)# clear counters`	Clears the counters associated with this N port.
Step 7	`switch(san-ext-nport)# end` `switch#`	Exits the SAN extension tuner submode.

Verifying the SAN Extension Tuner Configuration

The **show** commands display the current SAN extension tuner settings for the Cisco MDS switch (see Examples 41-1 to 41-6).

Example 41-1 Displays Entries in the FLOGI Database

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE    VSAN    FCID        PORT NAME                 NODE NAME
--------------------------------------------------------------------------------
iscsi3/4     200     0x050000    12:00:00:00:00:00:00:56     10:00:00:00:00:00:00:00
```

Example 41-2 Displays Details for a VSAN Entry in the FLOGI Database

```
switch# show fcns database vsan 200
VSAN 200
---------------------------------------------------------------------------
FCID          TYPE    PWWN (VENDOR)             FC4-TYPE:FEATURE
---------------------------------------------------------------------------
0x020000       N       22:22:22:22:22:22:22:22    scsi-fcp
0x050000       N       12:00:00:00:00:00:00:56    scsi-fcp
```

Example 41-3 Displays All Virtual N Ports Configured on the Specified Interface

```
switch# show san-ext-tuner interface gigabitethernet 3/4 nport pWWN
12:00:00:00:00:00:00:56 vsan 200 counters
Statistics for nport
Node name 10:00:00:00:00:00:00:00 Port name 12:00:00:00:00:00:00:56
I/Os per second               : 148
   Read                       : 0%
   Write                      : 100%
Ingress MB per second         : 0.02 MBs/sec (Max -0.02 MBs/sec)
Egress  MB per second         : 73.97 MBs/sec (Max -75.47 MBs/sec))
Average Response time per I/O : Read - 0 us, Write - 13432 us
Maximum Response time per I/O : Read - 0 us, Write - 6953 us
Minimum Response time per I/O : Read - 0 us, Write - 19752 us
Errors                        : 0
```

Example 41-4 Displays N Ports Configured on a Specified Gigabit Ethernet Interface

```
switch# show san-ext-tuner interface gigabitethernet 3/1
--------------------------------------------------------------------------------
Interface           NODE NAME                 PORT NAME                 VSAN
--------------------------------------------------------------------------------
GigabitEthernet3/1  10:00:00:00:00:00:00:00  10:00:00:00:00:00:00:01  91
```

Example 41-5 Displays the Transfer Ready Size Configured for a Specified N Port

```
switch# show san-ext-tuner interface gigabitethernet 3/1 nport pWWN 10:0:0:0:0:0:0:1 vsan
91
Node name            : 10:00:00:00:00:00:00:00
Port name            : 10:00:00:00:00:00:00:01
Transfer ready size  : all
```

Example 41-6 Displays All Virtual N Ports Configured in This Switch

```
switch# show san-ext-tuner nports
--------------------------------------------------------------------------------
Interface           NODE NAME                PORT NAME                VSAN
--------------------------------------------------------------------------------
GigabitEthernet3/1  10:00:00:00:00:00:00:00  10:00:00:00:00:00:00:01  91
```

Default Settings

Table 41-1 lists the default settings for tuning parameters.

Table 41-1 Default Tuning Parameters

Parameters	Default
Tuning	Disabled.
Transfer ready size	Same as the transfer size in the SCSI **write** command.
Outstanding I/Os	1.
Number of transactions	1.
Data generation format	All-zero format.
filemark-frequency	0.

CHAPTER 42

Configuring iSCSI

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

Note The iSCSI feature is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

The Cisco MDS 9216i switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

Note For information on configuring Gigabit Ethernet interfaces, see Chapter 45, "Configuring IPv4 for Gigabit Ethernet Interfaces."

This chapter includes the following sections:

About iSCSI

Note The iSCSI feature is not supported on the Cisco Fabric Switch for HP c-Class Bladesystem and Cisco Fabric Switch for IBM BladeCenter.

The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see Figure 42-1).

Figure 42-1 Transporting iSCSI Requests and Responses for Transparent iSCSI Routing

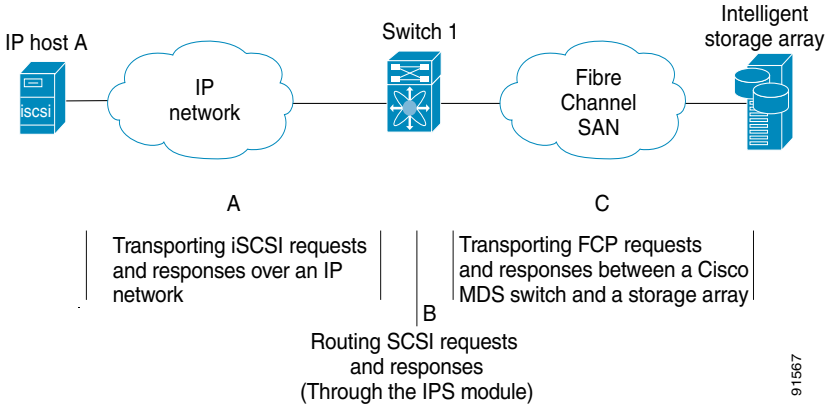

Each iSCSI host that requires access to storage through the IPS module or MPS-14/2 module needs to have a compatible iSCSI driver installed. (The Cisco.com website at http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420-scsi provides a list of compatible drivers.) Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver in the host.

The IPS module or MPS-14/2 module provides transparent SCSI routing. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. Figure 42-1 provides an example of a typical configuration of iSCSI hosts connected to an IPS module or MPS-14/2 module through the IP network access Fibre Channel storage on the Fibre Channel SAN.

The IPS module or MPS-14/2 module create a separate iSCSI SAN view and Fibre Channel SAN view. For the iSCSI SAN view, the IPS module or MPS-14/2 module creates iSCSI virtual targets and then maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical iSCSI targets were attached to the IP network (see Figure 42-2).

Figure 42-2 iSCSI SAN View—iSCSI virtual targets

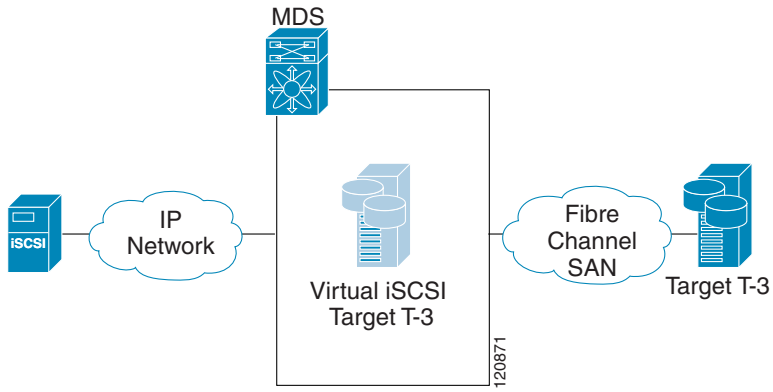

For the Fibre Channel SAN view, the IPS module or MPS-14/2 module presents iSCSI hosts as a virtual Fibre Channel host. The storage devices communicate with the virtual Fibre Channel host similar to communications performed with real Fibre Channel hosts (see Figure 42-3).

Figure 42-3 Fibre Channel SAN View—iSCSHI Host as an HBA

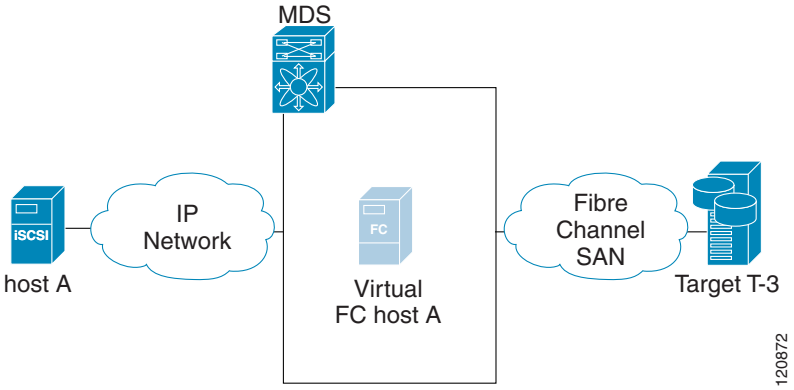

The IPS modules or MPS-14/2 modules transparently map the command between the iSCSI virtual target and the virtual Fibre Channel host (see Figure 42-4).

Figure 42-4 iSCSI to FCP (Fibre Channel) Routing

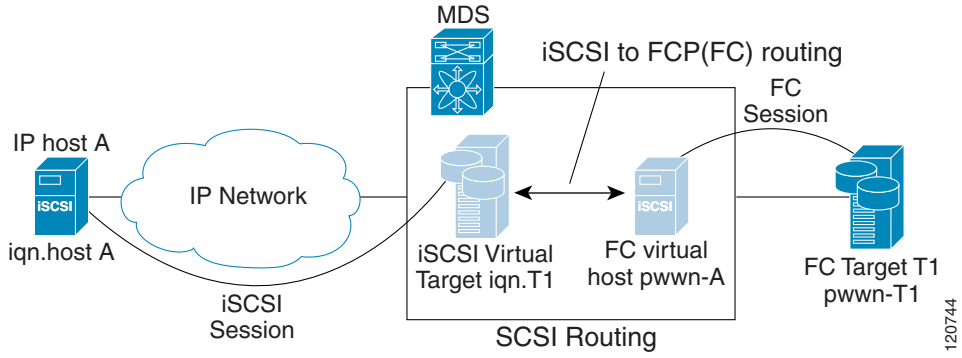

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions:

- The iSCSI requests and responses are transported over an IP network between the hosts and the IPS module or MPS-14/2 module.
- The SCSI requests and responses are routed between the hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). The IPS module or MPS-14/2 module performs this conversion and routing.
- The FCP requests or responses are transported between the IPS module or MPS-14/2 module and the Fibre Channel storage devices.

Note FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN. Refer to the IETF standards for IP storage at http://www.ietf.org for information on the iSCSI protocol.

About iSCSI Configuration Limits

iSCSI configuration has the following limits:

- The maximum number of iSCSI and iSLB initiators supported in a fabric is 2000.
- The maximum number of iSCSI and iSLB sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSCSI and iSLB session support by switch is 5000.
- The maximum number of iSCSI and iSLB targets supported in a fabric is 6000.

Configuring iSCSI

This section describes how to configure iSCSI on the Cisco MDS 9000 Family switches.

This section includes the following sections:

Enabling iSCSI

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

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

	Command	Purpose
Step 1	`switch# config t`	Enter configuration commands, one per line. End with CNTL/Z.
Step 2	`switch(config)# iscsi enable`	Enables iSCSI on that switch.
	`switch(config)# iscsi enable module <x>`	Enables iSCSI modules on the switch. **Note** New command added so that SME and iSCSI are available on the same switch.
	`switch(config)# no iscsi enable module <x>`	Disables the iSCSI module on the switch.
	`switch(config)# no iscsi enable`	Disables (default) iSCSI on that switch.

Caution When you disable this feature, all related configurations are automatically discarded.

Creating iSCSI Interfaces

Each physical Gigabit Ethernet interface on an IPS module or MPS-14/2 module can be used to translate and route iSCSI requests to Fibre Channel targets and responses in the opposite direction. To enable this capability, the corresponding iSCSI interface must be in an enabled state.

To enable iSCSI interfaces, follow these steps:

Step 1 Enable the required Gigabit Ethernet interface.

```
switch# config terminal
switch(config)# interface gigabitethernet 2/1
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

Step 2 Create the required iSCSI interface and enable the interface.

```
switch(config)# interface iscsi 2/1
switch(config-if)# no shutdown
```

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module or MPS-14/2 module presents physical Fibre Channel targets as iSCSI virtual targets, allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- Dynamic mapping—Automatically maps all the Fibre Channel target devices/ports as iSCSI devices. Use this mapping to create automatic iSCSI target names.
- Static mapping—Manually creates iSCSI target devices and maps them to the whole Fibre Channel target port or a subset of Fibre Channel LUNs. With this mapping, you must specify unique iSCSI target names.

  Static mapping should be used when iSCSI hosts should be restricted to subsets of LUs in the Fibre Channel targets and/or iSCSI access control is needed (see the "iSCSI Access Control" section on page 42-20). Also, static mapping allows the configuration of transparent failover if the LUs of the Fibre Channel targets are reachable by redundant Fibre Channel ports (see the "Transparent Target Failover" section on page 42-61).

Note The IPS module or MPS-14/2 module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module or MPS-14/2 module makes Fibre Channel targets available to iSCSI initiators.

Dynamic Mapping

When you configure dynamic mapping the IPS module or MPS-14/2 module imports all Fibre Channel targets to the iSCSI domain and maps each physical Fibre Channel target port as one iSCSI target. That is, all LUs accessible through the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the physical Fibre Channel target port.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

The IPS module or MPS-14/2 module creates an IQN formatted iSCSI target node name using the following conventions because the name must be unique in the SAN:

- IPS Gigabit Ethernet ports that are not part of a Virtual Router Redundancy Protocol (VRRP) group or PortChannel use this format:

  ```
  iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
  ```

- IPS ports that are part of a VRRP group use this format:

  ```
  iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
  ```

- Ports that are part of a PortChannel use this format:

  ```
  iqn.1987-02.com.cisco:02.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
  ```

Note If you have configured a switch name, then the switch name is used instead of the management IP address. If you have not configured a switch name, the management IP address is used.

With this convention, each IPS port in a Cisco MDS 9000 Family switch creates a unique iSCSI target node name for the same Fibre Channel target port in the SAN.

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0, LUN 1, and LUN 2, those LUNs would become available to an IP host through the iSCSI target node name iqn.1987-05.com.cisco:05.*MDS_switch_management_IP_address*.01-01.3100112233445566 (see Figure 42-5).

Figure 42-5 Dynamic Target Mapping

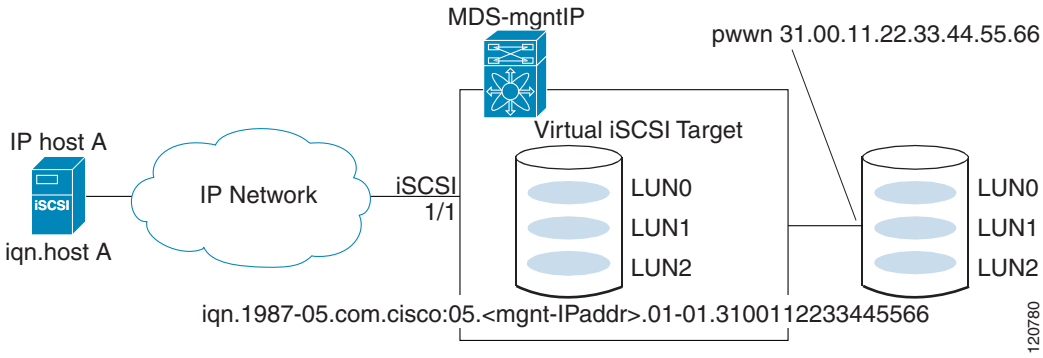

Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms (see the "iSCSI Access Control" section on page 42-20).

To enable dynamic mapping of Fibre Channel targets into iSCSI, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi import target fc`	IPS modules and MPS-14/2 modules dynamically import all Fibre Channel targets in the Fibre Channel SAN into the IP network.

Static Mapping

You can manually (statically) create an iSCSI target by assigning a user-defined unique iSCSI node name to it. The iSCSI qualified name is restricted to a minimum length of 16 characters and a maximum of 223 characters. A statically mapped iSCSI target can either map the whole Fibre Channel target port (all LUNs in the target port mapped to the iSCSI target), or it can contain one or more LUs from a Fibre Channel target port (see Figure 42-6).

Figure 42-6 Statically Mapped iSCSI Targets

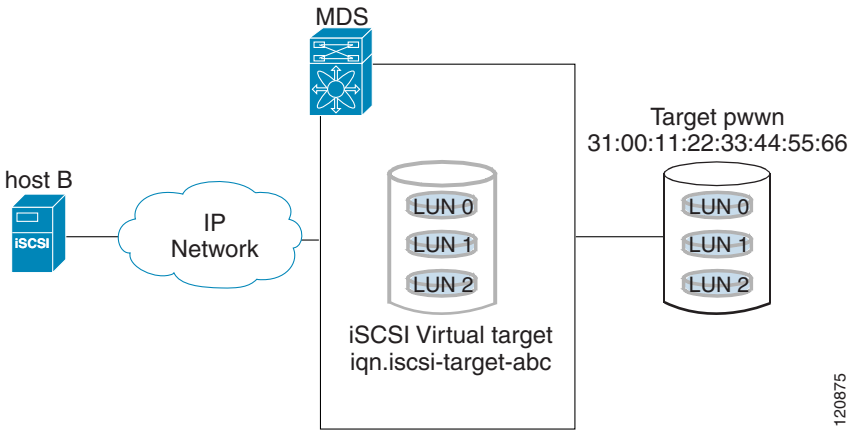

Advertising Static iSCSI Targets

You can limit the Gigabit Ethernet interfaces through which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To configure a specific interface that should advertise the iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	`switch(config-iscsi-tgt)# advertise interface GigabitEthernet 2/5`	Advertises the virtual target only on the specified interface. By default, it is advertised on all interfaces in all IPS modules or MPS-14/2 modules. **Note** To advertise the virtual target on multiple interfaces, issue the command for each interface.
	`switch(config-iscsi-tgt)# no advertise interface GigabitEthernet 2/5`	Removes this interface from the list of interfaces from which this target is advertised.

iSCSI Virtual Target Configuration Examples

This section provides three examples of iSCSI virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as an iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see Figure 42-7).

Figure 42-7 Assigning iSCSI Node Names

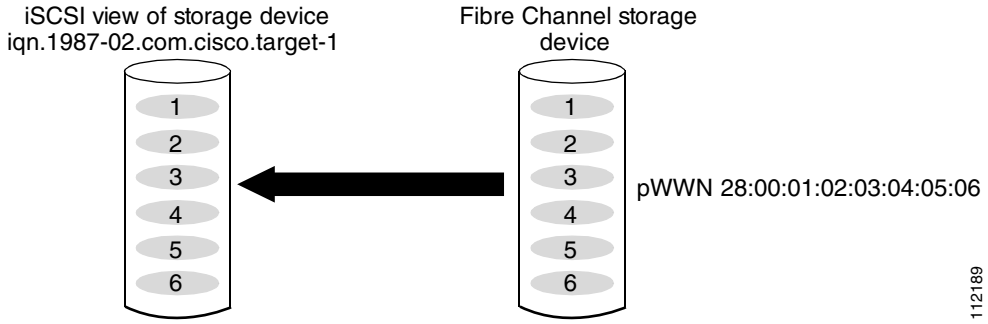


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06
```

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see Figure 42-8).

Figure 42-8 Mapping LUNs to an iSCSI Node Name

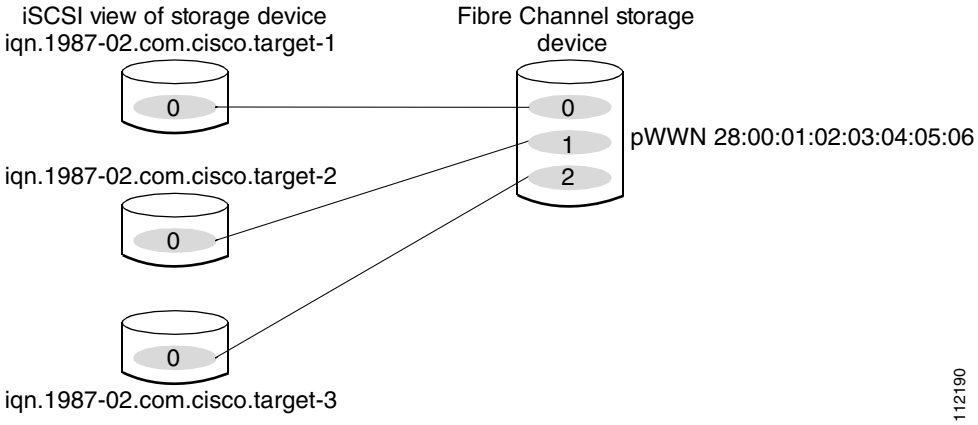


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
    pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
    pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```

Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see Figure 42-9).

Figure 42-9 Mapping LUNs to Multiple iSCSI Node Names

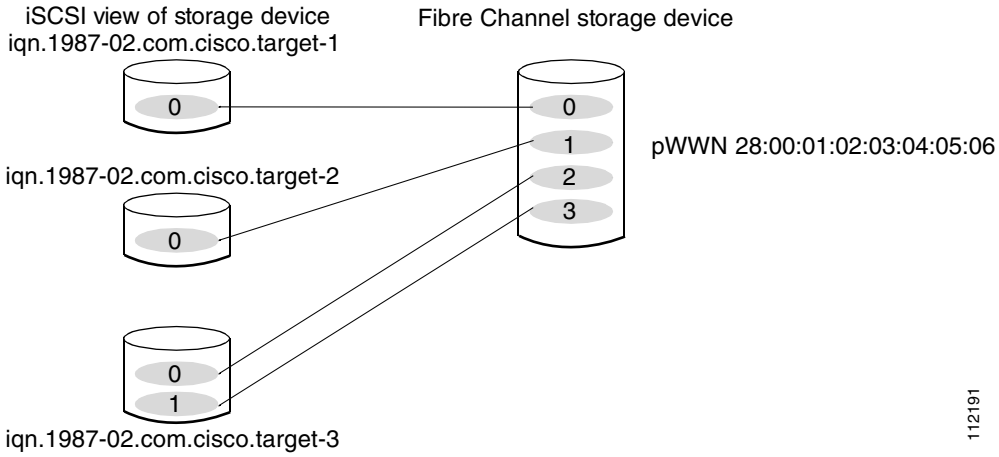


```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
    pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
    pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
    pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
    pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
```

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The IPS module or MPS-14/2 module connects to the Fibre Channel storage devices on behalf of the iSCSI host to send commands and transfer data to and from the storage devices. These modules use a virtual Fibre Channel N port to access the Fibre Channel storage devices on behalf of the iSCSI host. iSCSI hosts are identified by either iSCSI qualified name (IQN) or IP address.

Initiator Identification

iSCSI hosts can be identified by the IPS module or MPS-14/2 module using the following:

- iSCSI qualified name (IQN)

  An iSCSI initiator is identified based on the iSCSI node name it provides in the iSCSI login. This mode can be useful if an iSCSI host has multiple IP addresses and you want to provide the same service independent of the IP address used by the host. An initiator with multiple IP addresses (multiple network interface cards—NICs) has one virtual N port on each IPS port to which it logs in.

- IP address

  An iSCSI initiator is identified based on the IP address of the iSCSI host. This mode is useful if an iSCSI host has multiple IP addresses and you want to provide different service-based on the IP address used by the host. It is also easier to get the IP address of a host compared to getting the iSCSI

node name. A virtual N port is created for each IP address it uses to log in to iSCSI targets. If the host using one IP address logs in to multiple IPS ports, each IPS port will create one virtual N port for that IP address.

You can configure the iSCSI initiator identification mode on each IPS port and all the iSCSI hosts terminating on the IPS port will be identified according to that configuration. The default mode is to identify the initiator by name.

To specify the initiator identification mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface iscsi 4/1` `switch(config-if)#`	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	`switch(config-if)# switchport initiator id ip-address`	Identifies the iSCSI initiator based on the IP address.
	`switch(config-if)# switchport initiator id name`	Identifies the iSCSI initiator based on the initiator node name. This is the default behavior.

Initiator Presentation Modes

Two modes are available to present iSCSI hosts in the Fibre Channel fabric: transparent initiator mode and proxy initiator mode.

- In transparent initiator mode, each iSCSI host is presented as one virtual Fibre Channel host. The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration (similar to managing a “real” Fibre Channel host). Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.
- In -proxy initiator mode, there is only one virtual Fibre Channel host per one IPS port and all iSCSI hosts use that to access Fibre Channel targets. In a scenario where the Fibre Channel storage device requires explicit LUN access control for every host, the static configuration for each iSCSI initiator can be overwhelming. In such case, using the proxy- initiator mode simplifies the configuration.

Caution Enabling proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the “Changing iSCSI Interface Parameters and the Impact on Load Balancing” section on page 42-53.

The Cisco MDS switches support the following iSCSI session limits:

- The maximum number of iSCSI sessions on a switch is 5000.
- The maximum number of iSCSI sessions per IPS port in transparent initiator mode is 500.
- The maximum number of iSCSI sessions per IPS port in proxy initiator mode is 500.
- The maximum number of concurrent sessions an IPS port can create is five (but the total number of sessions that can be supported is 500).

Note If more than five iSCSI sessions try to come up simultaneously on a port, the initiator receives a temporary error and later retries to create a session.

Transparent Initiator Mode

Each iSCSI host is presented as one virtual Fibre Channel host (that is, one Fibre Channel N port). The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration. Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.

When an iSCSI host connects to the IPS module or MPS-14/2 module, a virtual host N port (HBA port) is created for the host (see Figure 42-10). Every Fibre Channel N port requires a unique Node WWN and Port WWN.

Figure 42-10 Virtual Host HBA Port

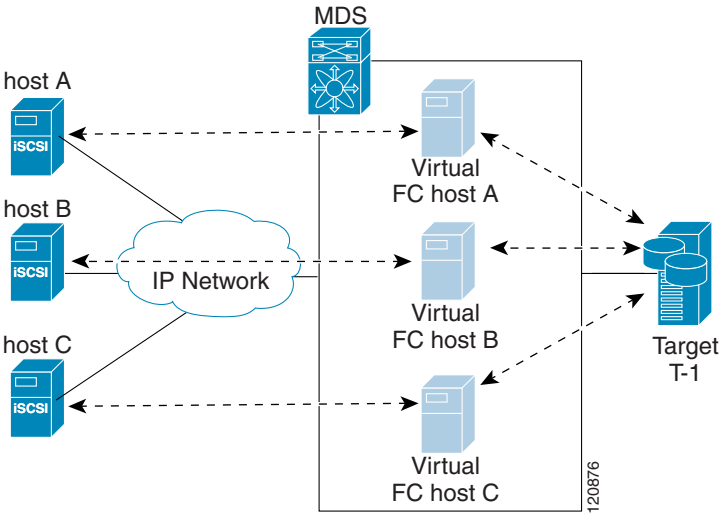

After the virtual N port is created with the WWNs, a fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the virtual N port is online in the Fibre Channel SAN and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- IP address of the iSCSI host in the IP-address field on the name server
- IQN of the iSCSI host in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor-specific iSCSI GW flag in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server.

When all the iSCSI sessions from the iSCSI host are terminated, the IPS modules or MPS-14/2 modules perform an explicit Fabric logout (FLOGO) to remove the virtual N-port device from the Fibre Channel SAN (this indirectly de-registers the device from the Fibre Channel name server).

For every iSCSI session from the host to the iSCSI virtual target there is a corresponding Fibre Channel session to the real Fibre Channel target. In Figure 42-10, there are three iSCSI hosts and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

iSCSI Initiator Idle Timeout

iSCSI initiator idle timeout specifies the time for which the virtual Fibre Channel N port is kept idle after the initiator logs out from its last iSCSI session. The default value for this timer is 300 seconds. This is useful to avoid N ports logging in to and logging off of the Fibre Channel SAN as transient failure occurs in the IP network. This helps reduce unnecessary RSCNs being generated in the Fibre Channel SAN.

To configure the initiator idle timeout, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator idle-timeout 10`	Configures the iSCSI initiators to have an idle timeout value of 10 seconds.

WWN Assignment for iSCSI Initiators

An iSCSI host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

Dynamic Mapping

With dynamic mapping, an iSCSI host is mapped to a dynamically generated port WWN (pWWN) and node WWN (nWWN). Each time the iSCSI host connects it might be mapped to a different WWN. Use this option if no access control is required on the Fibre Channel target device (because the target device access control is usually configured using the host WWN).

The WWNs are allocated from the MDS switch's WWN pool. The WWN mapping to the iSCSI host is maintained as long as the iSCSI host has at least one iSCSI session to the IPS port. When all iSCSI sessions from the host are terminated and the IPS module or MPS-14/2 module performs an FLOGO for the virtual N port of the host, the WWNs are released back to the switch's Fibre Channel WWN pool. These addresses are then available for assignment to other iSCSI hosts requiring access to the Fibre Channel Fabric.

The following are three dynamic initiator modes are supported:

- iSCSI—Dynamic initiators are treated as iSCSI initiators and can access dynamic virtual targets and configured iSCSI virtual targets.
- iSLB—Dynamic initiators are treated as iSLB initiators.
- Deny—Dynamic initiators are not allowed to log in to the MDS switch.

iSCSI dynamic mapping is the default mode of operation. This configuration is distributed using CFS.

Note Configuring dynamic initiator modes is supported only through the CLI, not through Device Manager or Fabric Manager.

To configure dynamic mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi dynamic initiator islb`	Specifies iSLB dynamic initiator mode.
	`switch(config)# iscsi dynamic initiator deny`	Disallows dynamic initiators from logging on to the MDS switch.
	`switch(config)# no iscsi dynamic initiator islb`	Reverts to iSCSI mode (default).

Static Mapping

With static mapping, an iSCSI host is mapped to a specific pWWN and nWWN. This mapping is maintained in persistent storage and each time the iSCSI host connects, the same WWN mapping is used. This mode is required if you use access control on the target device.

You can implement static mapping in one of two ways:

- User assignment—You can specify your own unique WWN by providing them during the configuration process.
- System assignment—You can request that the switch provide a WWN from the switch's Fibre Channel WWN pool and keep the mapping in its configuration.

Tip We recommend using the **system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the "World Wide Names" section on page 29-14). You should not use any previously assigned WWNs.

To configure static mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-init)#`	Configures an iSCSI initiator using the iSCSI name of the initiator node. The maximum name length is restricted to 223 alphanumeric characters. The minimum length is 16.
	`switch(config)# no iscsi initiator name iqn.1987-02.com.cisco.initiator`	Deletes the configured iSCSI initiator.

To configure static mapping (using the **ip-address** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator ip-address 10.50.0.0` `switch(config-iscsi-init)#`	Configures an iSCSI initiator using the IPv4 address of the initiator node.
	`switch(config)# iscsi initiator ip-address 2001:0DB8:800:200C::417A` `switch(config-iscsi-init)#`	Configures an iSCSI initiator using the IPv6 unicast address of the initiator node.
	`switch(config)# no iscsi initiator ip-address 2001:0DB8:800:200C::417A`	Deletes the configured iSCSI initiator.

To assign the WWN for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	`switch(config-iscsi-init)# static nWWN system-assign`	Uses the switch's WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent.
	`switch(config-iscsi-init)# static nWWN 20:00:00:05:30:00:59:11`	Assigns the user provided WWN as the nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node.
Step 2	`switch(config-iscsi-init)# static pWWN system-assign 2`	Uses the switch's WWN pool to allocate two pWWNs for this iSCSI initiator and keeps them persistent. The range is from 1 to 64.
	`switch(config-iscsi-init)# static pWWN 21:00:00:20:37:73:3b:20`	Assigns the user provided WWN as the pWWN for the iSCSI initiator.

Note If the system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is saved to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a write erase, you must manually delete the WWN configuration from the ASCII file. Failing to do so can cause duplicate WWN assignments if the ASCII configuration file is reapplied on the switch.

Making the Dynamic iSCSI Initiator WWN Mapping Static

After a dynamic iSCSI initiator has already logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping so this initiator uses the same mapping the next time it logs in.

You can convert a dynamic iSCSI initiator to static iSCSI initiator and make its WWNs persistent (see "Dynamic Mapping" section on page 42-13).

Note You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2s	`switch(config)# iscsi save-initiator name iqn.1987-02.com.cisco.initiator`	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose name is specified.
	`switch(config)# iscsi save-initiator ip-address 10.10.100.11`	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv4 address is specified.
	`switch(config)# iscsi save-initiator ip-address 2001:0DB8:800:200C::417A`	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv6 unicast address is specified.
	`switch(config)# iscsi save-initiator`	Saves the nWWN and pWWNs that have automatically been assigned to all the initiators.
Step 3	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 4	`switch# copy running-config startup-config`	Saves the nWWN/pWWN mapping configuration across system reboots.

Checking for WWN Conflicts

WWNs assigned to static iSCSI initiators by the system can be inadvertently returned to the system when an upgrade fails or you downgrade the system software (manually booting up an older Cisco MDS SAN-OS release without using the **install all** command). In these instances, the system can later assign those WWNs to other iSCSI initiators (dynamic or static) and cause conflicts.

You can address this problem by checking for and removing any configured WWNs that belong to the system whenever such scenarios occur.

To check for and remove WWN conflicts, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 1	`switch(config)# iscsi duplicate-wwn-check`
	`List of Potential WWN Conflicts:`
	`------------------------------`
	`Node : iqn.test-local-nwwn:1-local-pwwn:1`
	`nWWN : 22:03:00:0d:ec:02:cb:02`
	`pWWN : 22:04:00:0d:ec:02:cb:02`
		Checks for WWN conflicts.

Step 2s	`switch(config)# iscsi initiator name iqn.test-local-nwwn:1-local-pwwn:1`	Enters iSCSI initiator configuration mode for the initiator named iqn.test-local-nwwn:1-local-pwwn:1.
Step 3	`switch(config-iscsi-init)# no static nWWN 22:03:00:0d:ec:02:cb:02`	Removes a conflicting nWWN.
Step 4	`switch(config-iscsi-init)# no static pWWN 22:04:00:0d:ec:02:cb:02`	Removes a conflicting pWWN.

Proxy- Initiator Mode

In the event that the Fibre Channel storage device requires explicit LUN access control for every host using the transparent initiator mode (presenting one iSCSI host as one Fibre Channel host) means every iSCSI host has to be configured statically. This can mean several configuration tasks for each iSCSI host. In this case, using the proxy initiator mode simplifies the configuration.

In this mode, only one virtual host N port (HBA port) is created per IPS port. All the iSCSI hosts connecting to that IPS port will be multiplexed using the same virtual host N port (see Figure 42-11). This mode simplifies the task of statically binding WWNs. LUN mapping and assignment on the Fibre Channel storage array must be configured to allow access from the proxy virtual N port's pWWN for all LUNs used by each iSCSI initiator that connects through this IPS port. The LUN is then assigned to each iSCSI initiator by configuring iSCSI virtual targets (see the "Static Mapping" section on page 42-8) with LUN mapping and iSCSI access control (see the "iSCSI Access Control" section on page 42-20).

Figure 42-11 Multiplexing IPS Ports

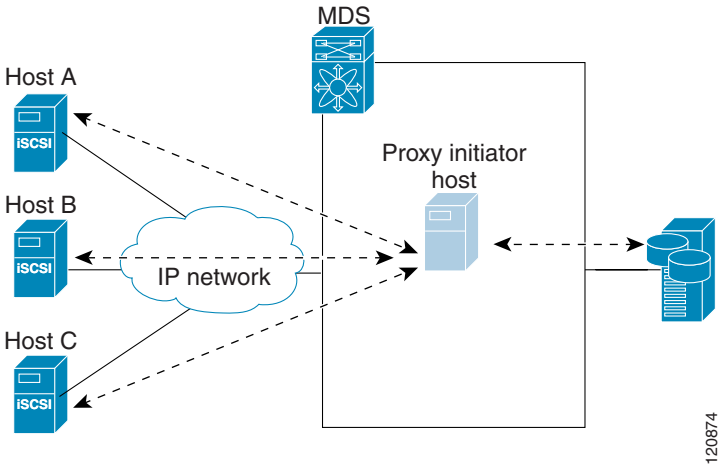

Proxy initiator mode can be configured on a per IPS port basis, in which case only iSCSI initiators terminating on that IPS port will be in this mode.

When an IPS port is configured in proxy-initiator mode, fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the proxy-initiator virtual N port is online in the Fibre Channel fabric and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- iSCSI interface name iSCSI slot /port is registered in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor specific flag (iscsi-gw) in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server

Similar to transparent initiator mode, the user can provide a pWWN and nWWN or request a system assigned WWN for the proxy initiator N port.

Caution Enabling the proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 42-53.

To configure the proxy initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface iscsi 4/1` `switch(config-if)#`	Selects the iSCSI interface on the switch that initiators will connect to.
Step 3	`switch(config-if)# switchport proxy-initiator`	Configures the proxy initiator mode with system-assignment nWWN and pWWN.
	`switch(config-if)# no switchport proxy-initiator`	Disables the proxy initiator mode.
Step 4	`switch(config-if)# switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22`	(Optional) Configures the proxy initiator mode using the specified WWNs.
	`switch(config-if)# no switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22`	Disables the proxy initiator mode.

Note When an interface is in proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the iSCSI interface's proxy N port attributes—the WWN pairs or the FC ID. You cannot configure zoning using iSCSI attributes such as IP address or IQN of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the "iSCSI Access Control" section on page 42-20).

VSAN Membership for iSCSI

Similar to Fibre Channel devices, iSCSI devices have two mechanisms by which VSAN membership can be defined.

- iSCSI host—VSAN membership to iSCSI host. (This method takes precedent over the iSCSI interface.)
- iSCSI interface—VSAN membership to iSCSI interface. (All iSCSI hosts connecting to this iSCSI interface inherit the interface VSAN membership if the host is not configured in any VSAN by the iSCSI host method.)

VSAN Membership for iSCSI Hosts

Individual iSCSI hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel, see Chapter 21, "Creating Dynamic VSANs"). The specified VSAN overrides the iSCSI interface VSAN membership.

To assign VSAN membership for iSCSI hosts, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-init)#`	Configures an iSCSI initiator.
Step 3	`switch(config-iscsi-init)# vsan 3`	Assigns the iSCSI initiator node to a specified VSAN. **Note** You can assign this host to one or more VSANs.
	`switch(config-iscsi-init)# no vsan 5`	Removes the iSCSI node from the specified VSAN.

Note When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

VSAN Membership for iSCSI Interfaces

VSAN membership can be configured for an iSCSI interface, called the *port VSAN*. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. In other words, the port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1.

Caution Changing the VSAN membership of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 42-53.

To change the default port VSAN for an iSCSI interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi interface vsan-membership`	Enables you to configure VSAN membership for iSCSI interfaces.
Step 3	`switch(config)# vsan database` `switch(config-vsan-db)#`	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.
Step 4	`switch(config-vsan-db)# vsan 2 interface iscsi 2/1`	Assigns the membership of the iscsi 2/1 interface to the specified VSAN (VSAN 2).
	`switch(config-vsan-db)# no vsan 2 interface iscsi 2/1`	Reverts to using the default VSAN as the port VSAN of the iSCSI interface.

Example of VSAN Membership for iSCSI Devices

Figure 42-12 provides an example of VSAN membership for iSCSI devices:

- iSCSI interface 1/1 is a member of VSAN Y.
- iSCSI initiator host A has explicit VSAN membership to VSAN X.
- Three iSCSI initiators (host A, host B, and host C) C connect to iSCSI interface 1/1.

Figure 42-12 VSAN Membership for iSCSI Interfaces

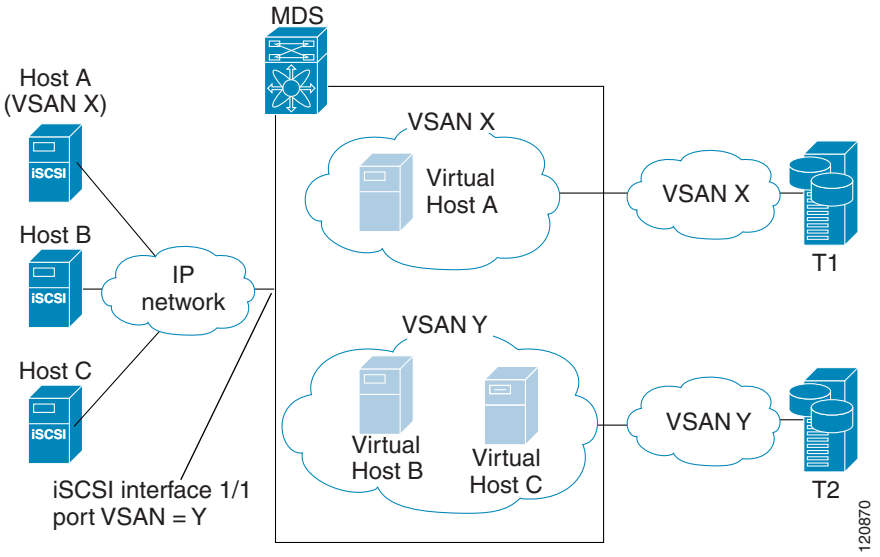

Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so they will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

Advanced VSAN Membership for iSCSI Hosts

An iSCSI host can be a member of multiple VSANs. In this case multiple virtual Fibre Channel hosts are created, one in each VSAN in which the iSCSI host is a member. This configuration is useful when certain resources such as Fibre Channel tape devices need to be shared among different VSANs.

iSCSI Access Control

Two mechanisms of access control are available for iSCSI devices.

- Fibre Channel zoning-based access control
- iSCSI ACL-based access control

Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel fabric, either or both the access control mechanisms can be used.

The topics included in this section are:

- Fibre Channel Zoning-Based Access Control, page 42-21

- iSCSI-Based Access Control, page 42-22
- Enforcing Access Control, page 42-23

Fibre Channel Zoning-Based Access Control

Cisco SAN-OS VSAN and zoning concepts have been extended to cover both Fibre Channel devices and iSCSI devices. Zoning is the standard access control mechanism for Fibre Channel devices, which is applied within the context of a VSAN. Fibre Channel zoning has been extended to support iSCSI devices, and this extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN.

Common mechanisms for identifying members in a Fibre Channel zone are the following (see Chapter 23, "Configuring and Managing Zones" for details on Fibre Channel zoning):

- Fibre Channel device pWWN.
- Interface and switch WWN. Device connecting via that interface is within the zone.

In the case of iSCSI, behind an iSCSI interface multiple iSCSI devices may be connected. Interface-based zoning may not be useful because all the iSCSI devices behind the interface will automatically be within the same zone.

In transparent initiator mode (where one Fibre Channel virtual N port is created for each iSCSI host as described in the "Transparent Initiator Mode" section on page 42-12), if an iSCSI host has static WWN mapping then the standard Fibre Channel device pWWN-based zoning membership mechanism can be used.

Zoning membership mechanism has been enhanced to add iSCSI devices to zones based on the following:

- IPv4 address/subnet mask
- IPv6 address/prefix length
- iSCSI qualified name (IQN)
- Symbolic-node-name (IQN)

For iSCSI hosts that do not have a static WWN mapping, the feature allows the IP address or iSCSI node name to be specified as zone members. Note that iSCSI hosts that have static WWN mapping can also use these features. IP address based zone membership allows multiple devices to be specified in one command by providing the subnet mask.

Note In proxy initiator mode, all iSCSI devices connecting to an IPS port gain access to the Fibre Channel fabric through a single virtual Fibre Channel N port. Thus, zoning based on the iSCSI node name or IP address will not have any effect. If zoning based on pWWN is used, then all iSCSI devices connecting to that IPS port will be put in the same zone. To implement individual initiator access control in proxy initiator mode, configure an iSCSI ACL on the virtual target (see the "iSCSI-Based Access Control" section on page 42-22).

To add an iSCSI initiator to the zone database, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# zone name iSCSIzone vsan 1` `switch(config-zone)`	Creates a zone name for the iSCSI devices in the IPS module or MPS-14/2 module to be included.
Step 3	`switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1`	Assigns an iSCSI node name-based membership into a zone.
	`switch(config-zone)# no member symbolic-nodename iqn.1987-02.com.cisco.init1`	Deletes the specified device from a zone.
	`switch(config-zone)# member ip-address 10.50.1.1`	Assigns an iSCSI IPv4 address-based membership into a zone.
	`switch(config-zone)# no member ip-address 10.50.1.1`	Deletes the specified device from a zone.
	`switch(config-zone)# member ipv6-address 2001:0DB8:800:200C::417A`	Assigns an iSCSI IPv6 address-based membership into a zone.
	`switch(config-zone)# no member ipv6-address 2001:0DB8:800:200C::417A`	Deletes the specified device from a zone.
	`switch(config-zone)# member pwwn 20:00:00:05:30:00:59:11`	Assigns an iSCSI port WWN-based membership into a zone.
	`switch(config-zone)# no member pwwn 20:00:00:05:30:00:59:11`	Deletes the device identified by the port WWN from a zone.

iSCSI-Based Access Control

iSCSI-based access control is applicable only if static iSCSI virtual targets are created (see the "Static Mapping" section on page 42-8). For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets.

By default, static iSCSI virtual targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow an iSCSI virtual target to be accessed by all hosts. The initiator access list can contain one or more initiators. The iSCSI initiator can be identified by one of the following mechanisms:

- iSCSI node name
- IPv4 address and subnet
- IPv6 address

Note For a transparent mode iSCSI initiator, if both Fibre Channel zoning and iSCSI ACLs are used, then for every static iSCSI target that is accessible to the iSCSI host, the initiator's virtual N port should be in the same Fibre Channel zone as the Fibre Channel target.

To configure access control in iSCSI, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-tgt)#`	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.

	Command	Purpose
Step 3	`switch(config-iscsi-tgt)# pWWN 26:00:01:02:03:04:05:06` `switch(config-iscsi-tgt)#`	Maps a virtual target node to a Fibre Channel target.
Step 4	`switch(config-iscsi-tgt)# initiator iqn.1987-02.com.cisco.initiator1 permit`	Allows the specified iSCSI initiator node to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	`switch(config-iscsi-tgt)# no initiator iqn.1987-02.com.cisco.initiator1 permit`	Prevents the specified initiator node from accessing virtual targets.
	`switch(config-iscsi-tgt)# initiator ip address 10.50.1.1 permit`	Allows the specified IPv4 address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	`switch(config-iscsi-tgt)# no initiator ip address 10.50.1.1 permit`	Prevents the specified IPv4 address from accessing virtual targets.
	`switch(config-iscsi-tgt)# initiator ip address 10.50.1.0 255.255.255.0 permit`	Allows all initiators in this IPv4 subnetwork (10.50.1/24) to access this virtual target.
	`switch(config-iscsi-tgt)# no initiator ip address 10.50.1.0 255.255.255.0 permit`	Prevents all initiators in this IPv4 subnetwork from accessing virtual targets.
	`switch(config-iscsi-tgt)# initiator ip address 2001:0DB8:800:200C::417A permit`	Allows the specified IPv6 unicast address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	`switch(config-iscsi-tgt)# no initiator ip address 2001:0DB8:800:200C::417A permit`	Prevents the specified IPv6 address from accessing virtual targets.
	`switch(config-iscsi-tgt)# initiator ip address 2001:0DB8:800:200C::/64 permit`	Allows all initiators in this IPv6 subnetwork (2001:0DB8:800:200C::/64) to access this virtual target.
	`switch(config-iscsi-tgt)# no initiator ip address 2001:0DB8:800:200C::/64 permit`	Prevents all initiators in this IPv6 subnetwork from accessing virtual targets.
	`switch(config-iscsi-tgt)# all-initiator-permit`	Allows all initiator nodes to access this virtual target.
	`switch(config-iscsi-tgt)# no all-initiator-permit`	Prevents any initiator from accessing virtual targets (default).

Enforcing Access Control

IPS modules and MPS-14/2 modules use both iSCSI and Fibre Channel zoning-based access control lists to enforce access control. Access control is enforced both during the iSCSI discovery phase and the iSCSI session creation phase. Access control enforcement is not required during the I/O phase because the IPS module or MPS-14/2 module is responsible for the routing of iSCSI traffic to Fibre Channel.

- iSCSI discovery phase—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module or MPS-14/2 module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section. The IPS module or MPS-14/2 module does this by querying the Fibre Channel name server for all the devices in the same zone as the initiator in all VSANs. It then filters out the devices that are initiators by looking at the FC4-feature field of the FCNS entry. (If a device does not register as either initiator or target in the FC4-feature field, the IPS module or MPS-14/2 module will advertise it.) It then

responds to the iSCSI host with the list of targets. Each will have either a static iSCSI target name that you configure or a dynamic iSCSI target name that the IPS module or MPS-14/2 module creates for it (see the "Dynamic Mapping" section on page 42-6).

- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module or MPS-14/2 module verifies if the specified iSCSI target (in the session login request) is allowed by both the access control mechanisms described in the "iSCSI-Based Access Control" section on page 42-22.

  If the iSCSI target is a static mapped target, the IPS module or MPS-14/2 module verifies if the iSCSI host is allowed within the access list of the iSCSI target. If the IP host does not have access, its login is rejected. If the iSCSI host is allowed, it validates if the virtual Fibre Channel N port used by the iSCSI host and the Fibre Channel target mapped to the static iSCSI virtual target are in the same Fibre Channel zone.

  If the iSCSI target is an autogenerated iSCSI target, then the IPS module or MPS-14/2 module extracts the WWN of the Fibre Channel target from the iSCSI target name and verifies if the initiator and the Fibre Channel target is in the same Fibre Channel zone or not. If they are, then access is allowed.

The IPS module or MPS-14/2 module uses the Fibre Channel virtual N port of the iSCSI host and does a zone-enforced name server query for the Fibre Channel target WWN. If the FC ID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI Session Authentication

The IPS module or MPS-14/2 module supports the iSCSI authentication mechanism to authenticate the iSCSI hosts that request access to the storage devices. By default, the IPS modules or MPS-14/2 modules allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation, you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see Chapter 33, "Configuring RADIUS and TACACS+"). AAA authentication supports a RADIUS, TACACS+, or local authentication device.

The **aaa authentication iscsi** command enables AAA authentication for the iSCSI host and specifies the method to use.

To configure AAA authentication for an iSCSI user, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# aaa authentication iscsi default group RadServerGrp`	Uses RADIUS servers that are added in the group called RadServerGrp for the iSCSI CHAP authentication.
	`switch(config)# aaa authentication iscsi default group TacServerGrp`	Uses TACACS+ servers that are added in the group called TacServerGrp for the iSCSI CHAP authentication.
	`switch(config)# aaa authentication iscsi default local`	Uses the local password database for iSCSI CHAP authentication.

The sections included in this topic are:

Authentication Mechanism

You can configure iSCSI CHAP or None authentication at both the global level and at each interface level.

The authentication for a Gigabit Ethernet interface or subinterface overrides the authentication method configured at the global level.

If CHAP authentication is used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used at all, issue the **iscsi authentication none** command.

To configure the authentication mechanism for iSCSI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi authentication chap`	Configures CHAP as the default authentication mechanism globally for the Cisco MDS switch. CHAP authentication is required for all iSCSI sessions.

To configure the authentication mechanism for iSCSI sessions to a particular interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface GigabitEthernet 2/1.100` `switch(config-if)#`	Selects the Gigabit Ethernet interface.
Step 3	`switch(config-if)# iscsi authentication none`	Specifies that no authentication is required for iSCSI sessions to the selected interface.

Local Authentication

See the"Characteristics of Strong Passwords" section on page 31-12 to create the local password database. To create users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory.

To configure iSCSI users for local authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# username iscsiuser password ffsffsfsffs345353554535 iscsi`	Configures a user name (iscsiuser) and password (ffsffsfsffs345353554535) in the local database for iSCSI login authentication.

Restricting iSCSI Initiator Authentication

By default, the iSCSI initiator can use any user name in the RADIUS server or in the local database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSCSI initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password has been compromised.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator name iqn.1987-02.com.cisco.init` `switch(config-iscsi-init)#`	Enters the configuration submode for the initiator iqn.1987-02.com.cisco.init.
Step 3	`switch(config-iscsi-init)# username user1`	Restricts the initiator `iqn.1987-02.com.cisco.init` to only authenticate using `user1` as its CHAP user name. **Tip** Be sure to define user1 as an iSCSI user in the local AAA database or the RADIUS server.

Mutual CHAP Authentication

In addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator, the IPS module or MPS-14/2 module also supports a mechanism for the iSCSI initiator to authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSCSI initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

To configure a global iSCSI target user name and password to be used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi authentication username testuser password abc123`	Configures the switch user account (testuser) along with a password (abc123) specified in clear text (default) for all initiators. The password is limited to 128 characters.
	`switch(config)# iscsi authentication username user1 password 7 !@*asdsfsdfjh!@df`	Configures the switch user account (user1) along with the encrypted password specified by 7 (!@*asdsfsdfjh!@df) for all initiators.
	`switch(config)# iscsi authentication username user1 password 0 abcd12AAA`	Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (indicated by 0—default) for all initiators. The password is limited to 128 characters.
	`switch(config)# no iscsi authentication username testuser`	Removes the global configuration for all initiators.

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-init)#`	Configures an iSCSI initiator using the iSCSI name of the initiator node.
Step 3	`switch(config-iscsi-init)# mutual-chap username testuser password abcd12AAA`	Configures the switch user account (testuser) along with a password (abcd12AAA) specified in clear text (default). The password is limited to 128 characters.
	`switch(config-iscsi-init)# mutual-chap username user1 password 7 !@*asdsfsdfjh!@df`	Configures the switch user account (user1) along with the encrypted password specified by 7 (!@*asdsfsdfjh!@df).
	`switch(config-iscsi-init)# no mutual-chap username testuser`	Removes the switch authentication configuration.

Use the **show running-config** and the **show iscsi global** commands to display the global configuration. Use the **show running-config** and the **show iscsi initiator configured** commands to display the initiator specific configuration.(See the "Displaying iSCSI Information" section on page 42-31 for command output examples.)

iSCSI Immediate Data and Unsolicited Data Features

Cisco MDS switches support the iSCSI immediate data and unsolicited data features if requested by the initiator during the login negotiation phase. Immediate data is iSCSI write data contained in the data segment of an iSCSI command protocol data unit (PDU), such as combining the write command and write data together in one PDU. Unsolicited data is iSCSI write data that an initiator sends to the iSCSI target, such as an MDS switch, in an iSCSI data-out PDU without having to receive an explicit ready to transfer (R2T) PDU from the target.

These two features help reduce I/O time for small write commands because it removes one round-trip between the initiator and the target for the R2T PDU. As an iSCSI target, the MDS switch allows up to 64 KB of unsolicited data per command. This is controlled by the FirstBurstLength parameter during iSCSI login negotiation phase.

If an iSCSI initiator supports immediate data and unsolicited data features, these features are automatically enabled on the MDS switch with no configuration required.

iSCSI Interface Advanced Features

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. These configurations are similar to the advanced FCIP configurations and are already explained in that section (see the "Displaying FCIP Profile Configuration Information" section on page 40-16).

To access these commands from the iSCSI interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface iscsi 4/1` `switch(config-if)#`	Selects the iSCSI interface on the switch.

Cisco MDS switches support the following advanced features for iSCSI interfaces:

- iSCSI Listener Port, page 42-28
- TCP Tuning Parameters, page 42-28
- QoS, page 42-29
- iSCSI Routing Modes, page 42-29

iSCSI Listener Port

You can configure the TCP port number for the iSCSI interface that listens for new TCP connections. The default port number is 3260. Once you change the TCP port number, the iSCSI port only accepts TCP connections on the newly configured port.

See the "Configuring TCP Listener Ports" section on page 40-11.

TCP Tuning Parameters

You can configure the following TCP parameters:

- Minimum retransmit timeout (See the "Minimum Retransmit Timeout" section on page 40-12.)
- Keepalive timeout (See the "Keepalive Timeout" section on page 40-12.)
- Maximum retransmissions (See the "Maximum Retransmissions" section on page 40-13)
- Path MTU (See the "Path MTUs" section on page 40-13.)
- SACK (SACK is enabled by default for iSCSI TCP configurations.) (See the "Selective Acknowledgments" section on page 40-13.)
- Window management (The iSCSI defaults are max-bandwidth is 1 Gbps, min-available-bandwidth is 70 Mbps, and round-trip-time is 1 msec.) (See the "Window Management" section on page 40-14.)
- Buffer size (The iSCSI default send buffer size is 4096 KB) (See the "Buffer Size" section on page 40-16.)
- Window congestion monitoring (enabled by default and the default burst size is 50 KB) (See the "Monitoring Congestion" section on page 40-15.)
- Maximum delay jitter (enabled by default and the default time is 500 microseconds) (See the "Estimating Maximum Jitter" section on page 40-15.)

QoS

To set the QoS values, follow these steps:

	Command	Purpose
Step 1	`switch(config-if)# qos 3`	Configures the differentiated services code point (DSCP) value of 3 to be applied to all outgoing IP packets in this iSCSI interface. The valid range for the iSCSI DSCP value is from 0 to 63.
Step 2	`switch(config-if)# no qos 5`	Reverts the switch to its factory default (marks all packets with DSCP value 0).

iSCSI Routing Modes

Cisco MDS 9000 Family switches support multiple iSCSI routing modes. Each mode negotiates different operational parameters, has different advantages and disadvantages, and is suitable for different usages.

- Pass-thru mode

  In pass-thru mode, the port on the IPS module or MPS 14/2 module converts and forwards read data frames from the Fibre Channel target to the iSCSI host frame-by-frame without buffering. This means that one data-in frame received is immediately sent out as one iSCSI data-in PDU.

  In the opposite direction, the port on the IPS module or MPS 14/2 module limits the maximum size of iSCSI write data-out PDU that the iSCSI host can send to the maximum data size that the Fibre Channel target specifies that it can receive. The result is one iSCSI data-out PDU received sent out as one Fibre Channel data frame to the Fibre Channel target.

  The absence of buffering in both directions leads to an advantage of lower forwarding latency. However, a small maximum data segment length usually results in lower data transfer performance from the host because of a higher processing overhead by the host system. Another benefit of this mode is iSCSI data digest can be enabled. This helps protect the integrity of iSCSI data carried in the PDU over what TCP checksum offers.

- Store-and-forward mode (default)

  In store-and-forward mode, the port on the IPS module or MPS 14/2 module assembles all the Fibre Channel data frames of an exchange to build one large iSCSI data-in PDU before forwarding it to the iSCSI client.

  In the opposite direction, the port on the IPS module or MPS 14/2 module does not impose a small data segment size on the host so the iSCSI host can send an iSCSI data-out PDU of any size (up to 256 KB). The port then waits until the whole iSCSI data-out PDU is received before it converts, or splits, the PDU, and forwards Fibre Channel frames to the Fibre Channel target.

  The advantage of this mode is higher data transfer performance from the host. The disadvantages are higher transfer latency and that the iSCSI data digest (CRC) cannot be used.

Note The store-and-forward mode is the default forwarding mode.

- Cut-through mode

  Cut-through mode improves the read operation performance over store-and-forward mode. The port on the IPS module or MPS 14/2 module achieves this by forwarding each Fibre Channel data-in frame to the iSCSI host as it is received without waiting for the whole exchange complete. There is no difference for write data-out operations from store-and-forward mode.

Figure 42-13 compares the messages exchanged by the iSCSI routing modes.

Figure 42-13 iSCSI Routing Modes

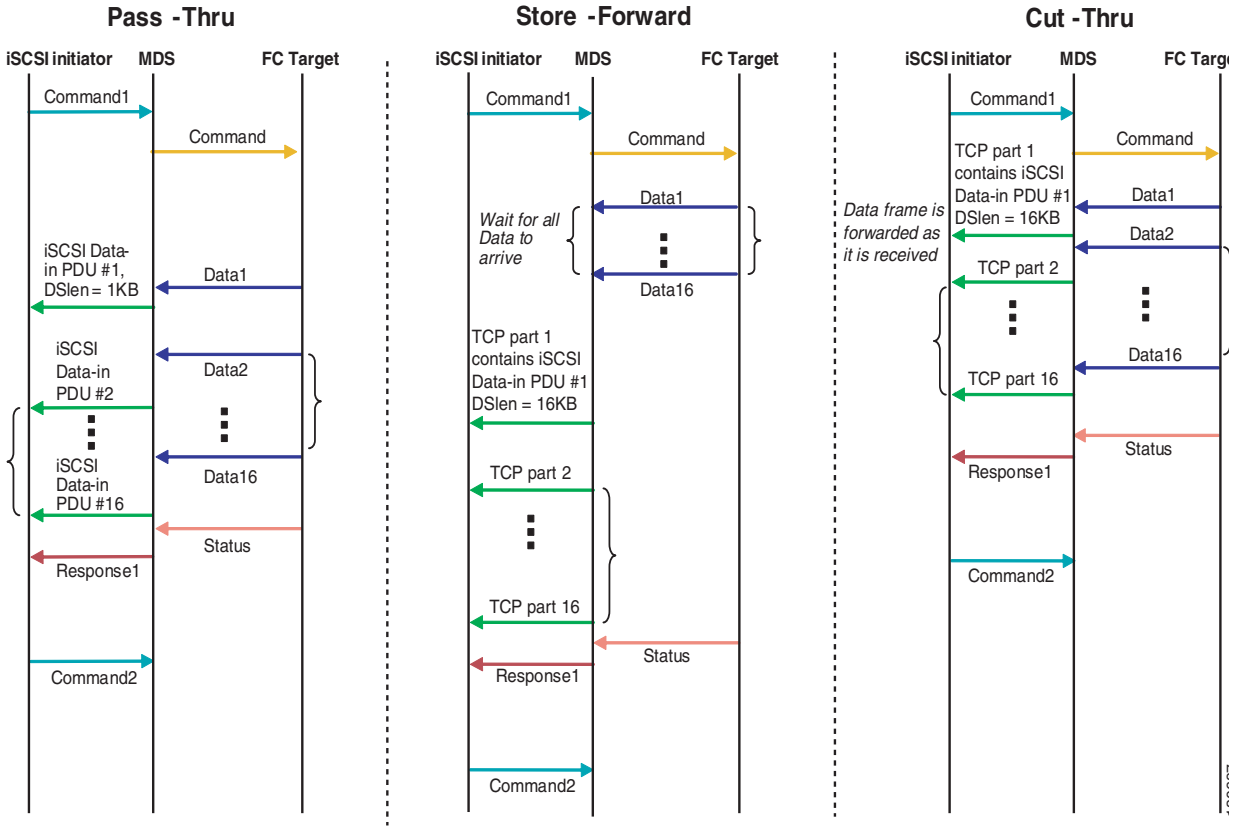

Table 42-1 compares the advantages and disadvantages of the different iSCSI routing modes.

Table 42-1 Comparison of iSCSI Routing Modes

Mode	Advantages	Disadvantages
Pass-thru	Low-latency Data digest can be used	Lower data transfer performance.
Store-and-forward	Higher data transfer performance	Data digest cannot be used.
Cut-thru	Improved read performance over store-and-forward	If the Fibre Channel target sent read data for different commands interchangeably, data of the first command is forwarded in cut-thru mode but the data of subsequent commands is buffered and the behavior is the same as store-and-forward mode. Data digest cannot be used.

Caution Changing the forwarding mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 42-53.

To set the iSCSI routing mode, follow this step:

	Command	Purpose
Step 1	`switch(config-if)# mode cut-thru`	Configures cut-thru mode on the iSCSI interface. **Caution** Changing the iSCSI routing mode disrupts the iSCSI sessions on the interface.
	`switch(config-if)# no mode cut-thru`	Reverts store-and-forward mode (default).

Displaying iSCSI Information

Use the **show iscsi** command to obtain detailed information about iSCSI configurations.

This section includes the following topics:

Displaying iSCSI Interfaces

Use the **show iscsi interface** command to view the summary, counter, description, and status of the iSCSI interface. Use the output to verify the administrative mode, the interface status, TCP parameters currently used, and brief statistics.

Example 42-1 Displays the iSCSI Interface Information

```
switch# show interface iscsi 4/1
iscsi4/1 is up
    Hardware is GigabitEthernet
    Port WWN is 20:cf:00:0c:85:90:3e:80
    Admin port mode is ISCSI
    Port mode is ISCSI
    Speed is 1 Gbps
    iSCSI initiator is identified by name
    Number of iSCSI session: 0 (discovery session: 0)
    Number of TCP connection: 0
    Configured TCP parameters
        Local Port is 3260
        PMTU discover is enabled, reset timeout is 3600 sec
        Keepalive-timeout is 60 sec
```

```
        Minimum-retransmit-time is 300 ms
        Max-retransmissions 4
        Sack is enabled
        QOS code point is 0
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 70000 kbps
        Estimated round trip time is 1000 usec
        Send buffer size is 4096 KB
        Congestion window monitoring is enabled, burst size is 50 KB
        Configured maximum jitter is 500 us
    Forwarding mode: store-and-forward
    TMF Queueing Mode : disabled
    Proxy Initiator Mode : disabled
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    iSCSI statistics
      Input 0 packets, 0 bytes
        Command 0 pdus, Data-out 0 pdus, 0 bytes
      Output 0 packets, 0 bytes
        Response 0 pdus (with sense 0), R2T 0 pdus
        Data-in 0 pdus, 0 bytes
```

Displaying iSCSI Statistics

Use the **show iscsi stats** command to view brief or detailed iSCSI statistics per iSCSI interface. See Example 42-2 and Example 42-3.

Example 42-2 displays iSCSI throughput on an IPS port in both inbound and outbound directions. It also displays the number of different types of iSCSI PDU received and transmitted by this IPS port.

Example 42-2 Display Brief iSCSI Statistics for an iSCSI Interface

```
switch# show iscsi stats iscsi 2/1
iscsi2/1
    5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
    5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
    iSCSI statistics
      974756 packets input, 142671620 bytes
        Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
bytes
      output 1022920 packets, 143446248 bytes
        Response 2352 pdus (with sense 266), R2T 1804 pdus
        Data-in 90453 pdus, 92458248 bytes
```

Example 42-3 displays detailed iSCSI statistics for an IPS port. Along with the traffic rate and the number of each iSCSI PDU type, it shows the number of FCP frames received and forwarded, the number of iSCSI login attempts, successes, and failures. It also shows the number of different types of iSCSI PDUs sent and received that are noncritical or occur less frequently, such as NOP in and out (NOP-In and NOP-Out), text request and response (Text-REQ and Text-RESP), and task management request and response (TMF-REQ and TMF-RESP).

Various types of errors and PDU or frame drop occurrences are also counted and displayed. For example, Bad header digest shows the number of iSCSI PDUs received that have a header digest that fails CRC verification. The iSCSI Drop section shows the number of PDUs that were dropped because of reasons such as target down, LUN mapping fail, Data CRC error, or unexpected Immediate or Unsolicited data. These statistics are helpful for debugging purposes when the feature is not working as expected.

The last section, Buffer Stats, gives statistics on the internal IPS packet buffer operation. This section is for debugging purposes only.

Example 42-3 Displays Detailed iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 2/1 detail
iscsi2/1
    5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
    5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
    iSCSI statistics
      974454 packets input, 142656516 bytes
        Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
bytes
      output 1022618 packets, 143431144 bytes
        Response 2352 pdus (with sense 266), R2T 1804 pdus
        Data-in 90453 pdus, 92458248 bytes
  iSCSI Forward:
    Command:2352 PDUs (Rcvd:2352)
    Data-Out (Write):16236 PDUs (Rcvd 44198), 0 fragments, 92364800 bytes, unsolicited 0
bytes
  FCP Forward:
    Xfer_rdy:1804 (Rcvd:1804)
    Data-In:90453 (Rcvd:90463), 92458248 bytes
    Response:2352 (Rcvd:2362), with sense 266
    TMF Resp:0

  iSCSI Stats:
    Login:attempt:13039, succeed:110, fail:12918, authen fail:0
    Rcvd:NOP-Out:914582, Sent:NOP-In:914582
          NOP-In:0, Sent:NOP-Out:0
          TMF-REQ:0, Sent:TMF-RESP:0
          Text-REQ:18, Sent:Text-RESP:27
          SNACK:0
          Unrecognized Opcode:0, Bad header digest:0
          Command in window but not next:0, exceed wait queue limit:0
          Received PDU in wrong phase:0
          SCSI Busy responses:0
    Immediate data failure::Separation:0
    Unsolicited data failure::Separation:0, Segment:0
                            Add header:0
    Sequence ID allocation failure:0
  FCP Stats:
    Total:Sent:47654
           Received:96625 (Error:0, Unknown:0)
    Sent:PLOGI:10, Rcvd:PLOGI_ACC:10, PLOGI_RJT:0
          PRLI:10, Rcvd:PRLI_ACC:10, PRLI_RJT:0, Error:0, From initiator:0
          LOGO:4, Rcvd:LOGO_ACC:0, LOGO_RJT:0
          PRLO:4, Rcvd:PRLO_ACC:0, PRLO_RJT:0
          ABTS:0, Rcvd:ABTS_ACC:0
          TMF REQ:0
          Self orig command:10, Rcvd:data:10, resp:10
    Rcvd:PLOGI:156, Sent:PLOGI_ACC:0, PLOGI_RJT:156
          LOGO:0, Sent:LOGO_ACC:0, LOGO_RJT:0
          PRLI:8, Sent:PRLI_ACC:8, PRLI_RJT:0
          PRLO:0, Sent:PRLO_ACC:0, PRLO_RJT:0
          ADISC:0, Sent:ADISC_ACC:0, ADISC_RJT:0
          ABTS:0

  iSCSI Drop:
    Command:Target down 0, Task in progress 0, LUN map fail 0
             CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
             No task:0
    Data-Out:0, Data CRC Error:0
    TMF-Req:0, No task:0
    Unsolicited data:0, Immediate command PDU:0
  FCP Drop:
    Xfer_rdy:0, Data-In:0, Response:0
```

```
  Buffer Stats:
    Buffer less than header size:0, Partial:45231, Split:322
    Pullup give new buf:0, Out of contiguous buf:0, Unaligned m_data:0
```

Displaying Proxy Initiator Information

If the proxy initiator feature is enabled in the iSCSI interface, use the **show interface iscsi** command to display configured proxy initiator information (see Example 42-4 and Example 42-5).

Example 42-4 Displays Proxy Initiator Information for the iSCSI Interface with System-Assigned WWNs

```
switch# show interface iscsi 4/1
iscsi4/1 is up
    Hardware is GigabitEthernet
    Port WWN is 20:c1:00:05:30:00:a7:9e
    Admin port mode is ISCSI
    Port mode is ISCSI
    Speed is 1 Gbps
    iSCSI initiator is identified by name
    Number of iSCSI session: 0, Number of TCP connection: 0
    Configured TCP parameters
        Local Port is 3260
        PMTU discover is enabled, reset timeout is 3600 sec
        Keepalive-timeout is 60 sec
        Minimum-retransmit-time is 300 ms
        Max-retransmissions 4
        Sack is disabled
        QOS code point is 0
    Forwarding mode: pass-thru
    TMF Queueing Mode : disabled
    Proxy Initiator Mode : enabled<----------------------------Proxy initiator is enabled
        nWWN is 28:00:00:05:30:00:a7:a1  (system-assigned)<----System-assigned nWWN
        pWWN is 28:01:00:05:30:00:a7:a1  (system-assigned)<---- System-assigned pWWN
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    iSCSI statistics
      Input 7 packets, 2912 bytes
        Command 0 pdus, Data-out 0 pdus, 0 bytes
      Output 7 packets, 336 bytes
        Response 0 pdus (with sense 0), R2T 0 pdus
        Data-in 0 pdus, 0 bytes
```

Example 42-5 Displays Proxy Initiator Information for the iSCSI Interface with User-Assigned WWNs

```
switch# show interface iscsi 4/2
iscsi4/2 is up
    Hardware is GigabitEthernet
    Port WWN is 20:c1:00:05:30:00:a7:9e
    Admin port mode is ISCSI
    Port mode is ISCSI
    Speed is 1 Gbps
    iSCSI initiator is identified by name
    Number of iSCSI session: 0, Number of TCP connection: 0
    Configured TCP parameters
        Local Port is 3260
        PMTU discover is enabled, reset timeout is 3600 sec
        Keepalive-timeout is 60 sec
```

```
        Minimum-retransmit-time is 300 ms
        Max-retransmissions 4
        Sack is disabled
        QOS code point is 0
    Forwarding mode: pass-thru
    TMF Queueing Mode : disabled
    Proxy Initiator Mode : enabled
        nWWN is 11:11:11:11:11:11:11:11 (manually-configured)<----User-assigned nWWN
        pWWN is 22:22:22:22:22:22:22:22 (manually-configured)<----User-assigned pWWN
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    iSCSI statistics
      Input 7 packets, 2912 bytes
        Command 0 pdus, Data-out 0 pdus, 0 bytes
      Output 7 packets, 336 bytes
        Response 0 pdus (with sense 0), R2T 0 pdus
        Data-in 0 pdus, 0 bytes
```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See Example 42-6.

Example 42-6 Displays the Current Global iSCSI Configuration and State

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP, NONE
  Import FC Target: Enabled
  Initiator idle timeout: 300 seconds
  Number of target node: 0
  Number of portals: 11
  Number of session: 0
  Failed session: 0, Last failed initiator name:
```

Displaying iSCSI Sessions

Use the **show iscsi session** command to view details about the current iSCSI sessions in the switch. Without parameters, this command displays all sessions. The output can be filtered by specifying an initiator, a target, or both.

Example 42-7 displays one iSCSI initiator configured based on the IQN (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on its IPv4 address (10.10.100.199).

Example 42-7 Displays Brief Information of All iSCSI Sessions

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  Session #1
    Discovery session, ISID 00023d000043, Status active

  Session #2
    Target VT1
    VSAN 1, ISID 00023d000046, Status active, no reservation
```

```
  Session #3
    Target VT2
    VSAN 1, ISID 00023d000048, Status active, no reservation

Initiator 10.10.100.199
  Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT2
    VSAN 1, ISID 246700000000, Status active, no reservation

  Session #2
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation

  Session #3
    Target iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
    VSAN 1, ISID 246e00000000, Status active, no reservation
```

Example 42-8 and Example 42-9 display the iSCSI initiator configured based on its IPv4 address (10.10.100.199).

Example 42-8 Displays Brief Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1
Initiator 10.10.100.199
  Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation
```

Example 42-9 Displays Detailed Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1 detail
Initiator 10.10.100.199 (oasis-qa)
  Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1 (index 3)
    Target VT1
    VSAN 1, ISID 246b00000000, TSIH 384, Status active, no reservation
    Type Normal, ExpCmdSN 39, MaxCmdSN 54, Barrier 0
    MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
    DataSeqInOrder No, InitialR2T Yes, ImmediateData No
    Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 38, Response: 38
      Bytes: TX: 8712, RX: 0
    Number of connection: 1
    Connection #1
      Local IP address: 10.10.100.200, Peer IP address: 10.10.100.199
      CID 0, State: LOGGED_IN
      StatSN 62, ExpStatSN 0
      MaxRecvDSLength 1024, our_MaxRecvDSLength 1392
      CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
      AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
      Version Min: 2, Max: 2
      FC target: Up, Reorder PDU: No, Marker send: No (int 0)
      Received MaxRecvDSLen key: No
```

Displaying iSCSI Initiators

Use the **show iscsi initiator** command to display information about all initiators connected to an iSCSI interface in the switch. The information can be filtered to display only the desired iSCSI initiator by specifying the initiator name. Detailed output of the iSCSI initiator can be obtained by specifying the **detail** option. The **iscsi-session** (and optionally **detail**) parameter displays only iSCSI session information. The **fcp-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See Example 42-10 and Example 42-11.

Example 42-10 Displays Information About Connected iSCSI Initiators

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
    Initiator ip addr (s): 10.10.100.116
    iSCSI alias name: AVANTI12-W2K
    Node WWN is 22:01:00:05:30:00:10:e1 (configured)
    Member of vsans: 1, 2, 10
    Number of Virtual n_ports: 1
    Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
      Interface iSCSI 4/1, Portal group tag: 0x180
      VSAN ID 1, FCID 0x6c0202
      VSAN ID 2, FCID 0x6e0000
      VSAN ID 10, FCID 0x790000

iSCSI Node name is 10.10.100.199
    iSCSI Initiator name: iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
    iSCSI alias name: oasis-qa
    Node WWN is 22:03:00:05:30:00:10:e1 (configured)
    Member of vsans: 1, 5
    Number of Virtual n_ports: 1
    Virtual Port WWN is 22:00:00:05:30:00:10:e1 (configured)
      Interface iSCSI 4/1, Portal group tag: 0x180
      VSAN ID 5, FCID 0x640000
      VSAN ID 1, FCID 0x6c0203
```

Example 42-11 Displays Detailed Information About the iSCSI Initiator

```
switch# show iscsi initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
    Initiator ip addr (s): 10.10.100.116
    iSCSI alias name: AVANTI12-W2K
    Node WWN is 22:01:00:05:30:00:10:e1 (configured)
    Member of vsans: 1, 2, 10
    Number of Virtual n_ports: 1

    Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
      Interface iSCSI 4/1, Portal group tag is 0x180
      VSAN ID 1, FCID 0x6c0202
      1 FC sessions, 1 iSCSI sessions
      iSCSI session details    <--------------------iSCSI session details
        Target: VT1
          Statistics:
            PDU: Command: 0, Response: 0
            Bytes: TX: 0, RX: 0
            Number of connection: 1
          TCP parameters
            Local 10.10.100.200:3260, Remote 10.10.100.116:4190
            Path MTU: 1500 bytes
            Retransmission timeout: 310 ms
            Round trip time: Smoothed 160 ms, Variance: 38
```

```
        Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
        Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
        Congestion window: Current: 1 KB

FCP Session details      <--------------------FCP session details
  Target FCID: 0x6c01e8 (S_ID of this session: 0x6c0202)
    pWWN: 21:00:00:20:37:62:c0:0c, nWWN: 20:00:00:20:37:62:c0:0c
    Session state: CLEANUP
    1 iSCSI sessions share this FC session
      Target: VT1
    Negotiated parameters
      RcvDataFieldSize 1392 our_RcvDataFieldSize 1392
      MaxBurstSize 0, EMPD: FALSE
      Random Relative Offset: FALSE, Sequence-in-order: Yes
    Statistics:
      PDU: Command: 0, Response: 0
```

Use the **show fcns database** (and optionally **detail**) to display the Fibre Channel name server entry for the Fibre Channel N port created for iSCSI initiators in the SAN. See Example 42-12 and Example 42-13.

Example 42-12 Displays the FCNS Database Contents

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
------------------------------------------------------------------------
0x020101    N     22:04:00:05:30:00:35:e1 (Cisco)         scsi-fcp:init isc..w <--iSCSI
0x020102    N     22:02:00:05:30:00:35:e1 (Cisco)         scsi-fcp:init isc..w initiator
0x0205d4    NL    21:00:00:04:cf:da:fe:c6 (Seagate)       scsi-fcp:target
0x0205d5    NL    21:00:00:04:cf:e6:e4:4b (Seagate)       scsi-fcp:target
...
Total number of entries = 10

VSAN 2:
------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
------------------------------------------------------------------------
0xef0001    N     22:02:00:05:30:00:35:e1 (Cisco)         scsi-fcp:init isc..w
Total number of entries = 1

VSAN 3:
------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
------------------------------------------------------------------------
0xed0001    N     22:02:00:05:30:00:35:e1 (Cisco)         scsi-fcp:init isc..w
Total number of entries = 1
```

Example 42-13 Displays the FCNS Database in Detail

```
switch# show fcns database detail
------------------------
VSAN:1     FCID:0x020101
------------------------
port-wwn (vendor)     :22:04:00:05:30:00:35:e1 (Cisco)
node-wwn              :22:03:00:05:30:00:35:e1
class                 :2,3
node-ip-addr          :10.2.2.12                 <--- iSCSI initiator's IPv4 address
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
```

```
symbolic-port-name    :
symbolic-node-name    :iqn.1991-05.com.microsoft:oasis2-dell <--- iSCSI initiator's IQN
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :22:01:00:05:30:00:35:de
hard-addr             :0x000000
------------------------
VSAN:1     FCID:0x020102
------------------------
port-wwn (vendor)     :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn              :22:01:00:05:30:00:35:e1
class                 :2,3
node-ip-addr          :10.2.2.11
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name    :
symbolic-node-name    :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :22:01:00:05:30:00:35:de
hard-addr             :0x000000
...
Total number of entries = 10
==========================================================================
------------------------
VSAN:2     FCID:0xef0001
------------------------
port-wwn (vendor)     :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn              :22:01:00:05:30:00:35:e1
class                 :2,3
node-ip-addr          :10.2.2.11
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name    :
symbolic-node-name    :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :22:01:00:05:30:00:35:de
hard-addr             :0x000000
Total number of entries = 1
...
```

Use the **show iscsi initiator configured** to display information about all the configured iSCSI initiators. Specifying the name shows information about the desired initiator. See Example 42-14.

Example 42-14 Displays Information About Configured Initiators

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
    Member of vsans: 1, 2, 10
    Node WWN is 22:01:00:05:30:00:10:e1
    No. of PWWN: 5
      Port WWN is 22:04:00:05:30:00:10:e1
      Port WWN is 22:05:00:05:30:00:10:e1
      Port WWN is 22:06:00:05:30:00:10:e1
      Port WWN is 22:07:00:05:30:00:10:e1
      Port WWN is 22:08:00:05:30:00:10:e1

iSCSI Node name is 10.10.100.199
    Member of vsans: 1, 5
    Node WWN is 22:03:00:05:30:00:10:e1
    No. of PWWN: 4
      Port WWN is 22:00:00:05:30:00:10:e1
```

```
        Port WWN is 22:09:00:05:30:00:10:e1
        Port WWN is 22:0a:00:05:30:00:10:e1
        Port WWN is 22:0b:00:05:30:00:10:e1

    User Name for Mutual CHAP: testuser
```

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the Fibre Channel targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See Example 42-15.

Example 42-15 Displays Exported Targets

```
switch# show iscsi virtual-target
target: VT1
  * Port WWN 21:00:00:20:37:62:c0:0c
    Configured node
    all initiator permit is enabled

target: VT2
    Port WWN 21:00:00:04:cf:4c:52:c1
    Configured node
    all initiator permit is disabled
target: iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
    Port WWN 21:00:00:20:37:a6:be:32 , VSAN 1
    Auto-created node
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See Example 42-16.

Example 42-16 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdffg

username:user2
secret:cshadhdhsadadjajdjas
```

Configuring iSLB

The iSCSI server load balancing (iSLB) feature provides a means to easily configure large scale iSCSI deployments containing hundreds or even thousands of initiators. When not using iSLB, configuring iSCSI requires the following:

- You need to perform multiple configuration steps on the MDS switch, including the following:
  - Initiator configuration using static pWWN and VSAN.
  - Zoning configuration for initiators and targets.
  - Optional create virtual target and give access to the initiator.
  - Configuration of target LUN mapping and masking on the storage system for the initiator based on the static pWWN created for the initiator on the MDS switch.
- You need to duplicate the configuration manually on multiple MDS switches.
- There is no load balancing for IPS ports. For example:
  - The Virtual Router Redundancy Protocol (VRRP) only supports active and backup, not load balancing.
  - You must use multiple VRRP groups and configure hosts in different groups.

iSLB provides the following features:

- The iSLB initiator configuration is simplified with support for initiator targets and auto-zones.
- Cisco Fabric Services (CFS) eliminates the need for manual configuration by distributing the iSLB initiator configuration among all MDS switches in the fabric.

Note Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically mapped iSCSI initiator configurations are not distributed.

- Dynamic load balancing of iSLB initiators is available using iSCSI login redirect and VRRP.

This section covers the following topics:

Note Before configuring iSLB, you must enable iSCSI (see the "Enabling iSCSI" section on page 42-5).

Note For iSLB, all switches in the fabric must be running Cisco MDS SAN-OS Release 2.1(1a) or later.

About iSLB Configuration Limits

iSLB configuration has the following limits:

- The maximum number of iSLB and iSCSI initiators supported in a fabric is 2000.
- The maximum number of iSLB and iSCSI sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSLBand iSCSI session support by switch is 5000.
- The maximum number of iSLB and iSCSI targets supported in a fabric is 6000.
- The maximum number of switches in a fabric that can have iSLB with CFS distribution enabled is four.
- No more than 200 new iSLB initiators can be added to the pending configuration. Before adding more initiators, you must commit the configuration.
- You cannot disable iSCSI if you have more than 200 iSLB initiators in the running configuration. Reduce the number of iSLB initiators to fewer than 200 before disabling iSCSI.
- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic is disrupted when any zoneset is activated.
- If IVR and iSLB features are enabled in the same fabric, you should have at least one switch in the fabric where both these features are enabled. Any zoning-related configuration and activation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, there may be traffic disruption in the fabric.

iSLB Configuration Prerequisites

Perform the following prerequisite actions prior to configuring iSLB:

- Enable iSCSI (see the "Enabling iSCSI" section on page 42-5).
- Configure the Gigabit Ethernet interfaces (see the "Basic Gigabit Ethernet Configuration for IPv4" section on page 45-2 or the Configuring Basic Connectivity for IPv6, page 46-11).
- Configure the VRRP groups (see the "Configuring Load Balancing Using VRRP" section on page 42-56).
- Configure and activate a zone set (see Chapter 23, "Configuring and Managing Zones").
- Enable CFS distribution for iSLB (see the "Enabling iSLB Configuration Distribution" section on page 42-58).

About iSLB Initiators

iSLB initiators provide the following features in addition to those supported by iSCSI initiators:

- An iSLB initiator also supports iSLB virtual targets. These targets are very similar to iSCSI virtual targets with the exception that they do not include the advertise interface option and as a result are distributable using CFS.
- Initiator targets—-These targets are configured for a particular initiator.
- Load balancing using iSCSI login redirect and VRRP—If load balancing is enabled, the IPS Manager redirects incoming sessions to the best interface based on the calculated load for each interface.
- Configuration distribution to other switches using CFS.

Configuring iSLB Initiators

This section includes the following topics:

Configuring iSLB Initiator Names or IP Addresses

You must specify the iSLB initiator name or IP address before configuring it.

Note Specifying the iSLB initiator name or IP address is the same as for an iSCSI initiator. See the "Static Mapping" section on page 42-14.

To enter iSLB initiator configuration submode using the **name** option for an iSLB initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator` `switch(config-islb-init)#`	Configures an iSLB initiator using the iSCSI name of the initiator node (iqn.1987-02.com.cisco.initiator) and enters iSLB initiator configuration submode. The maximum name length is 223 alphanumeric characters. The minimum length is 16.
	`switch(config)# no islb initiator name iqn.1987-02.com.cisco.initiator`	Deletes the configured iSLB initiator.

To enter iSLB initiator configuration submode using the **ip-address** option for an iSLB initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator ip-address 10.1.1.3` `switch(config-islb-init)#`	Configures an iSLB initiator using the IPv4 address of the initiator node and enters iSLB initiator configuration submode.
	`switch(config)# no islb initiator ip-address 10.1.1.3`	Deletes the configured iSLB initiator.
	`switch(config)# islb initiator ip-address 2001:0DB8:800:200C::417A` `switch(config-islb-init)#`	Configures an iSLB initiator using the IPv6 unicast address of the initiator node and enters iSLB initiator configuration submode.
	`switch(config)# no islb initiator ip-address 2001:0DB8:800:200C::417A`	Deletes the configured iSLB initiator.

Assigning WWNs to iSLB Initiators

An iSLB host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

Note Assigning WWNs for iSLB initiators is the same as for iSCSI initiators. For information on dynamic and static mapping, see the "WWN Assignment for iSCSI Initiators" section on page 42-13.

Tip We recommend using the **SystemAssign system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the "World Wide Names" section on page 29-14). You should not use any previously assigned WWNs.

Making the Dynamic iSLB Initiator WWN Mapping Static

After a dynamic iSLB initiator has logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping to allow this initiator to use the same mapping the next time it logs in .

You can convert a dynamic iSLB initiator to a static iSLB initiator and make its WWNs persistent (see "Dynamic Mapping" section on page 42-13).

Note You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.

Note Making the dynamic mapping for iSLB initiators static is the same as for iSCSI. See the "Making the Dynamic iSCSI Initiator WWN Mapping Static" section on page 42-15.

Note Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically configured iSCSI initiator configurations are not distributed.

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2s	`switch(config)# islb save-initiator name iqn.1987-02.com.cisco.initiator`	Saves the nWWNs and pWWNs that have automatically been assigned to the iSLB initiator whose name is specified.
	`switch(config)# islb save-initiator 10.10.100.11`	Saves the nWWNs and pWWNs that have automatically been assigned to the iSLB initiator whose IPv4 address is specified.
	`switch(config)# iscsi save-initiator ip-address 2001:0DB8:800:200C::417A`	Saves the nWWNs and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv6 unicast address is specified.
	`switch(config)# islb save-initiator`	Saves the nWWNs and pWWNs that have automatically been assigned to all the iSLB initiators.
Step 3	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 4	`switch# copy running-config startup-config`	Saves the nWWN/pWWN mapping configuration across system reboots.

Assigning VSAN Membership for iSLB Initiators

Individual iSLB hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel; see Chapter 21, "Creating Dynamic VSANs"). The specified VSAN overrides the iSCSI interface VSAN membership.

Note Specifying the iSLB initiator VSAN is the same as for an iSCSI initiator. See the "VSAN Membership for iSCSI" section on page 42-18.

To assign VSAN membership for iSLB initiators, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator ip-address 10.1.1.3` `switch(config-islb-init)#`	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submode.
Step 3	`switch(config-islb-init)# vsan 3`	Assigns the iSLB initiator node to a specified VSAN. **Note** You can assign this host to one or more VSANs.
	`switch(config-islb-init)# no vsan 3`	Removes the iSLB initiator from the specified VSAN.

Note When an iSLB initiator is configured in any other VSAN (other than VSAN 1, the default VSAN), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

Configuring Metric for Load Balancing

You can assign a load metric to each initiator for weighted load balancing. The load calculated is based on the number of initiators on a given iSCSI interface. This feature accommodates initiators with different bandwidth requirements. For example, you could assign a higher load metric to a a database server than to a web server. Weighted load balancing also accommodates initiators with different link speeds.

For more information on load balancing, see the "About Load Balancing Using VRRP" section on page 42-51.

To configure a weight for load balancing, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-init)#`	Configures an iSLB initiator using the name of the initiator node and enters iSLB initiator configuration mode.
Step 3	`switch(config-iscsi-init)# metric 100`	Assigns 100 as the weight metric for this iSLB initiator.
Step 4	`switch(config-iscsi-init)# no metric 100`	Reverts to the default value (1000).

Verifying iSLB Initiator Configuration

To verify the iSLB initiator configuration, use the **show islb initiator configured** command.

```
switch# show islb initiator configured
iSCSI Node name is 10.1.1.2
    Member of vsans: 10
    Node WWN is 23:02:00:0c:85:90:3e:82
    Load Balance Metric: 100
    Number of Initiator Targets: 1
```

```
Initiator Target: test-targt
  Port WWN 01:01:01:01:02:02:02:02
  Primary PWWN VSAN 1
  Zoning support is enabled
  Trespass support is disabled
  Revert to primary support is disabled
```

Configuring iSLB Initiator Targets

You can configure initiator targets using the device alias or the pWWN. You can also optionally specify one or more of the following optional parameters:

- Secondary pWWN
- Secondary device alias
- LUN mapping
- IQN
- VSAN identifier

Note The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.

In addition, you can disable auto-zoning.

If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

To configure iSLB initiator targets, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator ip-address 10.1.1.3` `switch(config-islb-init)#`	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submode.

	Command	Purpose
Step 3	`switch(config-iscsi-islb-init)# target pwwn 26:00:01:02:03:04:05:06`	Configures the iSLB initiator target using a pWWN with auto-zoning enabled (default).
	`switch(config-iscsi-islb-init)# target pwwn 26:00:01:02:03:04:05:06 no-zone`	Configures the iSLB initiator target using a pWWN with auto-zoning disabled.
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias`	Configures the iSLB initiator target using a device alias with auto-zoning enabled (default).
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias fc-lun 0x1234 iscsi-lun 0x2345`	Configures the iSLB initiator target using a device alias and optional LUN mapping. **Note** The CLI interprets the LUN identifier value as a hexadecimal value whether or not the **0x** prefix is included.
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias iqn-name iqn.1987-01.com.cisco.initiator`	Configures the iSLB initiator target using a device alias and an optional IQN.
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias sec-device-alias SecondaryAlias`	Configures the iSLB initiator target using a device alias and an optional secondary device alias.
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias sec-pwwn 26:01:02:03:04:05:06:07`	Configures the iSLB initiator target using a device alias and an optional secondary pWWN.
	`switch(config-iscsi-islb-init)# target device-alias SampleAlias vsan 10`	Configures the iSLB initiator target using a device alias and the VSAN identifier. **Note** The VSAN identifier is optional is if the target is online. If the target is not online, the VSAN identifier is required.
	`switch(config-iscsi-init)# no target pwwn 26:00:01:02:03:04:05:06`	Removes the iSLB initiator target.

Configuring and Activating Zones for iSLB Initiators and Initiator Targets

You can configure a zone name where the iSLB initiators and initiator targets are added. If you do not specify a zone name, the IPS manager creates one dynamically. iSLB zone sets have the following considerations:

- Auto-zoning of the initiator with the initiator targets is enabled by default.
- A zone set must be active in a VSAN for auto-zones to be created in that VSAN.
- iSLB zone set activation might fail if another zone set activation is in process or if the zoning database is locked. Retry the iSLB zone set activation if a failure occurs. To avoid this problem, only perform only one zoning related operation (normal zones, IVR zones, or iSLB zones) at a time.
- Auto-zones are created when the zone set is activated and there has been at least one change in the zoneset. The activation has no effect if only the auto-zones have changed.

Caution If IVR and iSLB are enabled in the same fabric, at least one switch in the fabric must have both features enabled. Any zoning related configuration or activation operation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, traffic might be disrupted in the fabric.

To configure the iSLB initiator optional auto-zone name and activate the zone set, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# islb initiator ip-address 10.1.1.3` `switch(config-islb-init)#`	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submode.
Step 3	`switch(config-islb-init)# zonename IslbZone`	Specifies the zone name where the initiators and the initiator targets are added (optional).
	`switch(config-islb-init)# no zonename IslbZone`	Removes the initiators and initiator targets from the zone and adds them to a dynamically created zone (default).
Step 4	`switch(config-islb-init)# exit`	Returns to configuration mode.
Step 5	`switch(config)# islb zoneset activate`	Activates zoning for the iSLB initiators and initiator targets with zoning enabled and creates auto-zones if no zone names are configured. **Note** This step is not required if CFS is enabled. CFS automatically activates the zone when the configuration changes are committed.

Verifying iSLB Zoning Configuration

The following example shows the **show zoneset active** command output when the dynamically generated zone name is used.

```
switch# show zoneset active
zoneset name zoneset-1 vsan 1
  zone name ips_zone_5d9603bcff68008a6fc5862a6670ca09 vsan 1
  * fcid 0x010009 [ip-address 10.1.1.3]
    pwwn 22:00:00:04:cf:75:28:4d
    pwwn 22:00:00:04:cf:75:ed:53
    pwwn 22:00:00:04:cf:75:21:d5
    pwwn 22:00:00:04:cf:75:ee:59
...
```

The following example shows the **show zoneset active** command output when the configured zone name IslbZone is used.

```
switch# show zoneset active
zoneset name zoneset-1 vsan 1
  zone name ips_zone_IslbZone vsan 1
    ip-address 10.1.1.3
    pwwn 22:00:00:04:cf:75:28:4d
    pwwn 22:00:00:04:cf:75:ed:53
    pwwn 22:00:00:04:cf:75:21:d5
    pwwn 22:00:00:04:cf:75:ee:59
...
```

Configuring iSLB Session Authentication

The IPS module and MPS-14/2 module support the iSLB authentication mechanism to authenticate iSLB hosts that request access to storage. By default, the IPS module and MPS-14/2 module allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see Chapter 33, "Configuring RADIUS and TACACS+"). AAA authentication supports RADIUS, TACACS+, or a local authentication device.

Note Specifying the iSLB session authentication is the same as for iSCSI. See the "iSCSI Session Authentication" section on page 42-24.

Restricting iSLB Initiator Authentication

By default, the iSLB initiator can use any user name in the RADIUS or local AAA database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSLB initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password have been compromised.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator name iqn.1987-02.com.cisco.init` `switch(config-islb-init)#`	Configures an iSLB initiator using the IQN of the initiator node and enters iSLB initiator configuration mode.
Step 3	`switch(config-islb-init)# username user1`	Restricts the initiator `iqn.1987-02.com.cisco.init` to only authenticate using `user1` as its CHAP user name. **Tip** Be sure to define user1 as an iSCSI user in the local AAA database or the RADIUS server.

Mutual CHAP Authentication

In addition to the IPS module and MPS-14/2 module authentication of the iSLB initiator, the IPS module and MPS-14/2 module also support a mechanism for the iSLB initiator to authenticate the Cisco MDS switch's initiator target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSLB initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

To configure a per-initiator user name and password used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator` `switch(config-islb-init)#`	Configures an iSLB initiator using the name of the initiator node and enters iSLB initiator configuration mode.

	Command	Purpose
Step 3	`switch(config-islb-init)# mutual-chap username testuser password dcba12LKJ`	Configures the switch user account (testuser) along with a password (dcba12LKJ) specified in clear text (default). The password is limited to 128 characters.
	`switch(config-islb-init)# mutual-chap username testuser password 7 !@*asdsfsdfjh!@df`	Configures the switch user account (testuser) along with the encrypted password specified by 7 (!@*asdsfsdfjh!@df).
Step 4	`switch(config-iscsi-init)# no mutual-chap username testuser`	Removes the switch authentication configuration.

Verifying iSLB Authentication Configuration

Use the **show running-config** and the **show iscsi global** (see Example 42-6) commands to display the global configuration. Use the **show running-config** and the **show islb initiator configured** (see Example 42-14) commands to display the initiator specific configuration.

To verify the iSLB user name and mutual CHAP configuration, use the **show islb initiator configured** command.

```
switch# show islb initiator configured
iSCSI Node name is 10.1.1.3
  Member of vsans: 3
  User Name for login authentication: user1
  User Name for Mutual CHAP: testuser
  Load Balance Metric: 1000  Number of Initiator Targets: 1
  Number of Initiator Targets: 1

  Initiator Target: iqn.1987-05.com.cisco:05.ips-hac4
    Port WWN 50:06:04:82:ca:e1:26:8d
    Zoning Enabled
    No. of LU mapping: 3
      iSCSI LUN: 0x0001, FC LUN: 0x0001
      iSCSI LUN: 0x0002, FC LUN: 0x0002
      iSCSI LUN: 0x0003, FC LUN: 0x0003
```

About Load Balancing Using VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB. Figure 42-14 shows an example of load balancing using iSLB.

Figure 42-14 iSLB Initiator Load Balancing Example

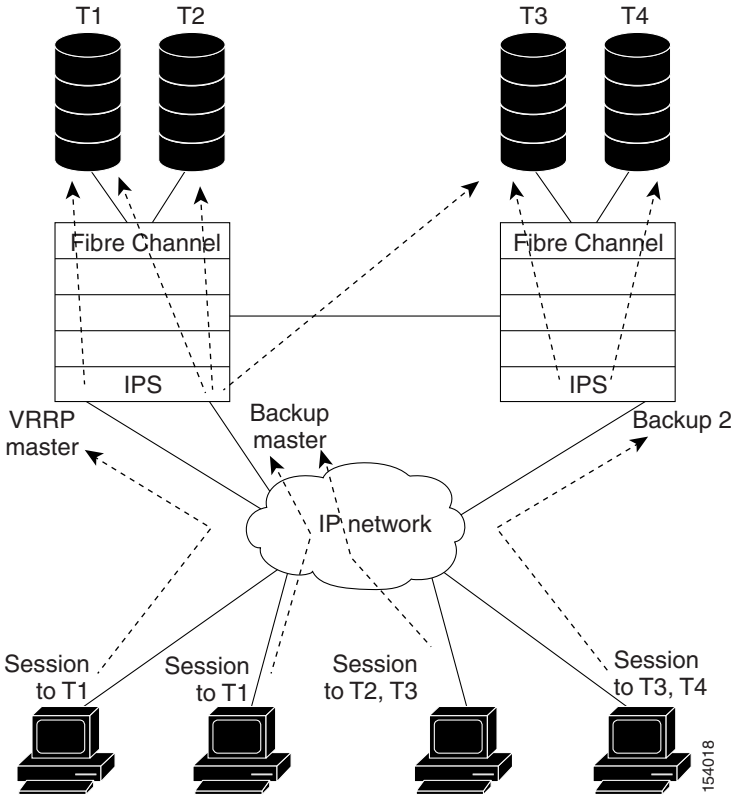

The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. This information is synchronized to all switches through CFS if recovery is needed when a master port fails. The initiator gets a temporary redirect iSCSI login response. The host then logs in to the backup port at its physical IP address. If the backup port goes down, the host will revert to the master port. The master port knows through CFS that the backup port has gone down and redirects the host to another backup port.

Note If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.

Note An initiator can also be redirected to the physical IP address of the master interface.

Tip iSLB VRRP load balancing is based on the number of iSLB initiators and not number of sessions. Any iSLB initiator that has more targets configured than the other iSLB initiators (resulting in more sessions) should be configured with a higher load metric. For example, you can increase the load metric of the iSLB initiator with more targets to 3000 from the default value of 1000.

Caution A Gigabit Ethernet interface configured for iSLB can only be in one VRRP group because redirected sessions do not carry information about the VRRP IP address or group. This restriction allows thebackup port to uniquely identify the VRRP group to which it belongs.

Changing iSCSI Interface Parameters and the Impact on Load Balancing

All iSCSI interfaces in a VRRP group that has load balancing enabled must have the same interface VSAN, authentication, proxy initiator mode, and forwarding mode. When you need to change any of these parameters for the iSCSI interfaces in a VRRP group, you must do so one interface at a time. During the transition time when the parameter is changed on some interfaces in the VRRP group and not the others, the master port does not redirect new initiators and instead handles them locally.

Caution Changing the VSAN, proxy initiator, authentication, and forwarding mode for iSCSI interfaces in a VRRP group can cause sessions to go down multiple times.

VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces

When the VRRP master receives an iSCSI session request from an initiator, it first checks for an existing mapping to one of the interfaces in that VRRP group. If such a mapping exists, the VRRP master redirects the initiator to that interface. If no such mapping exists, the VRRP master selects the least loaded interface and updates the selected interface's load with the initiator's iSLB metric (weight).

Note The VRRP master interface is treated specially and it takes lower load compared to the other interfaces. This is to account for the redirection work performed by the master interface for every session. A new initiator is assigned to the master interface only if the following is true for every other interface:

VRRP backup interface load > [2 * VRRP master interface load + 1]

The Example 42-17 and Example 42-18 are based the following configuration:

- GigabitEthernet2/1.441 is the VRRP master interface for Switch1.
- GigabitEthernet2/2.441 is the VRRP backup interface for Switch1.
- GigabitEthernet1/1.441 is the VRRP backup interface for Switch2.
- GigabitEthernet1/2.441 is the VRRP backup interface for Switch2.

Example 42-17 Load Distribution With the Default Metric

The follow example output shows the initial load distribution for three initiators with the default load metric value.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  0
  1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  1000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  1000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  1000
               -- Initiator To Interface Assignment --
--------------------------------------------------------------------------------
Initiator                VR Id VRRP IP      Switch WWN              Ifindex
--------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init  1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
```

The following example output shows load distribution for four initiators. The interface load metric value for the master interface changed from 0 to 1000.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VVR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  1000
  1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  1000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  1000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  1000
               -- Initiator To Interface Assignment --
--------------------------------------------------------------------------------
Initiator                VR Id VRRP IP      Switch WWN              Ifindex
--------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

The follow example output shows load distribution for nine initiators. The interface load metric values for the backup interfaces have changed.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VVR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  1000
  1     10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  3000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  3000
  1     10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  2000
                  -- Initiator To Interface Assignment --
--------------------------------------------------------------------------------
Initiator                VR Id VRRP IP      Switch WWN              Ifindex
--------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
iqn.cisco.test-linux.init4 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
```

```
iqn.cisco.test-linux.init5 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init6 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init7 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init8 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
```

Example 42-18 Load Distribution With the Metric Set to 3000 on One Initiator

The follow example output shows the initial load distribution for three initiators with one initiator having load metric of 3000 and the remaining initiator with the default metric value.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VVR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  0
  1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  1000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  3000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  1000
                    -- Initiator To Interface Assignment --
--------------------------------------------------------------------------------
Initiator               VR Id VRRP IP      Switch WWN              Ifindex
--------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init1 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init2 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
```

The follow example output shows load distribution for four initiators. The interface load metric value for the master interface changed from 0 to 1000.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VVR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  1000
  1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  3000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  1000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  1000
                    -- Initiator To Interface Assignment --
--------------------------------------------------------------------------------
Initiator               VR Id VRRP IP      Switch WWN              Ifindex
--------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init2 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init3 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

The following example output shows load distribution for nine initiators. The interface load metric values for the backup interfaces have changed.

```
switch# show islb vrrp summary
...
--------------------------------------------------------------------------------
VVR Id   VRRP IP        Switch WWN               Ifindex                 Load
--------------------------------------------------------------------------------
M 1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/1.441  2000
  1      10.10.122.115  20:00:00:0b:5f:3c:01:80   GigabitEthernet2/2.441  3000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/1.441  3000
  1      10.10.122.115  20:00:00:0c:ce:5c:5b:c0   GigabitEthernet1/2.441  3000
```

```
                    -- Initiator To Interface Assignment --
-------------------------------------------------------------------------------
Initiator                VR Id VRRP IP      Switch WWN              Ifindex
-------------------------------------------------------------------------------
iqn.cisco.test-linux.init0 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
iqn.cisco.test-linux.init4 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init5 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init6 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init7 1  10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init8 1  10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

Configuring Load Balancing Using VRRP

You must first configure VRRP on the Gigabit Ethernet interfaces on the switch that connect to the IP network before configuring VRRP for iSLB. For information on how to configure VRRP on a Gigabit Ethernet interface, see the "Virtual Router Redundancy Protocol" section on page 43-16.

Enabling VRRP for Load Balancing

To enable or disable VRRP for iSLB, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# islb vrrp 10 load-balance`	Enables iSLB VRRP for IPv4 VR group 10.
Step 3	`switch(config)# no islb vrrp 10 load-balance`	Disables iSLB VRRP for IPv4 VR group 10.
Step 4	`switch(config)# islb vrrp ipv6 20 load-balance`	Enables iSLB VRRP for IPv6 VR group 20.
Step 5	`switch(config)# no islb vrrp ipv6 20 load-balance`	Disables iSLB VRRP for IPv6 VR group 20.

Verifying iSLB VRRP Load Balancing Configuration

To verify the iSLB VRRP load balancing configuration for IPv4, use the **show vrrp vr** command.

```
switch# show vrrp vr 1
      Interface  VR IpVersion Pri   Time Pre State   VR IP addr
------------------------------------------------------------------------------
        GigE1/5   1   IPv4    100    1 s     master  10.10.10.1
        GigE1/6   1   IPv4    100    1 s     master  10.10.10.1
```

To verify the iSLB VRRP load balancing configuration for IPv6, use the **show vrrp ipv6 vr** command.

```
switch# show vrrp ipv6 vr 1
      Interface  VR IpVersion Pri   Time Pre State   VR IP addr
------------------------------------------------------------------------
        GigE6/2   1   IPv6    100  100cs     master  5000:1::100
       PortCh 4   1   IPv6    100  100cs     master  5000:1::100
```

Displaying iSLB VRRP Information

Use the **show islb vrrp summary vr** command to display VRRP load balancing information.

```
switch# show islb vrrp summary vr 30

                        -- Groups For Load Balance --
--------------------------------------------------------------------------------
VR Id             VRRP Address Type             Configured Status
--------------------------------------------------------------------------------
30                      IPv4                          Enabled

                       -- Interfaces For Load Balance --
--------------------------------------------------------------------------------
 VR Id      VRRP IP         Switch WWN                 Ifindex            Load
--------------------------------------------------------------------------------
  30   192.168.30.40 20:00:00:0d:ec:02:cb:00      GigabitEthernet3/1      2000
  30   192.168.30.40 20:00:00:0d:ec:02:cb:00      GigabitEthernet3/2      2000
  30   192.168.30.40 20:00:00:0d:ec:0c:6b:c0      GigabitEthernet4/1      2000
M 30   192.168.30.40 20:00:00:0d:ec:0c:6b:c0      GigabitEthernet4/2      1000
```

About iSLB Configuration Distribution Using CFS

Configuration for iSLB initiators and initiator targets on an MDS switch can be distributed using the Cisco Fabric Services (CFS). This feature allows you to synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, iSCSI dynamic initiator mode, and global authentication parameters are also distributed. CFS distribution is disabled by default (see Chapter 5, “Using the CFS Infrastructure”).

After enabling the distribution, the first configuration starts an implicit session. All server configuration changes entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database.

When CFS is enabled for iSLB, the first iSLB configuration operation starts a CFS session and locks the iSLB configuration in the fabric. The configuration changes are applied to the pending configuration database. When you make the changes to the fabric, the pending configuration is distributed to all the switches in the fabric. Each switch then validates the configuration. This check ensures the following:

- The VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names do not conflict with the iSCSI initiators on all the switches.

After the check completes successfully, all the switches commit the pending configuration to the running configuration. If any check fails, the entire commit fails.

Note iSLB is only fully supported when CFS is enabled. Using iSLB auto-zoning without enabling CFS mode may cause traffic disruption when any zone set is activated.

Note CFS does not distribute non-iSLB initiator configurations or import Fibre Channel target settings.

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

Distributing the iSLB Configuration Using CFS

This section contains the following:

Enabling iSLB Configuration Distribution

To enable CFS distribution of the iSLB configuration, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# islb distribute`	Enables iSLB configuration distribution.
	`switch(config)# no islb distribute`	Disables (default) iSLB configuration distribution.

Locking the Fabric

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. Once 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. Modifications from this point on 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.

Note iSCSI configuration changes are not allowed when an iSLB CFS session is active.

Committing Changes to the Fabric

To apply the pending iSLB configuration changes to the active configuration and to other MDS switches in the fabric, 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 MDS switches throughout the fabric, the automatic zones are activated, and the fabric lock is released.

To commit iSLB configuration changes to other MDS switches in the fabric, activate iSLB automatic zones, and release the fabric lock, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# islb commit`	Commits the iSLB configuration distribution, activates iSLB automatic zones, and releases the fabric lock.

Discarding Pending Changes

At any time, you can discard the pending changes to the iSLB configuration and release the fabric lock. This action has no affect on the active configuration on any switch in the fabric.

To discard the pending iSLB configuration changes and release the fabric lock, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# islb abort`	Commits the iSLB configuration distribution.

Clearing a Fabric Lock

If you have performed an iSLB 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.

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

To release a fabric lock, issue the **clear islb session** command in EXEC mode using a login ID that has administrative privileges.

```
switch# clear islb session
```

CFS Merge Process

When two fabrics merge, CFS attempts to merge the iSLB configuration from both the fabrics. A designated switch (called the *dominant switch*) in one fabric sends its iSLB configuration to a designated switch (called the *subordinate switch*) in the other fabric. The subordinate switch compares its running

configuration to the received configuration for any conflicts. If no conflicts are detected, it merges the two configurations and sends it to all the switches in both the fabrics. Each switch then validates the configuration. This check ensures the following:

- VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names have no conflicts with iSCSI initiators on all the switches.

If this check completes successfully, the subordinate switch directs all the switches to commit the merged configuration to running configuration. If any check fails, the merge fails.

The **show islb merge status** command displays the exact reason for the failure. The first successful commit request after a merge failure takes the fabric out of the merge failure state.

Displaying Pending iSLB Configuration Changes

You can display the pending configuration changes using the **show islb pending** command.

```
switch# show islb pending
iscsi initiator idle-timeout 10
islb initiator ip-address 10.1.1.1
static pWWN 23:01:00:0c:85:90:3e:82
static pWWN 23:06:00:0c:85:90:3e:82
username test1
islb initiator ip-address 10.1.1.2
static nWWN 23:02:00:0c:85:90:3e:82
```

You can display the differences between the pending configuration and the current configuration using the **show islb pending-diff** command.

```
switch# show islb pending-diff
+iscsi initiator idle-timeout 10
islb initiator ip-address 10.1.1.1
+ static pWWN 23:06:00:0c:85:90:3e:82
+islb initiator ip-address 10.1.1.2
+ static nWWN 23:02:00:0c:85:90:3e:82
```

Displaying iSLB CFS Status

You can display the iSLB CFS status using the **show islb session status** command.

```
switch# show islb status
iSLB Distribute is enabled
iSLB CFS Session exists
```

Displaying iSLB CFS Distribution Session Status

You can display the status of the iSLB CFS distribution session using the **show islb cfs-session status** command.

```
switch# show islb cfs-session status
last action               : fabric distribute enable
last action result        : success
last action failure cause : success
```

Displaying iSLB CFS Merge Status

You can display the iSLB CFS merge status using the **show islb merge status** command.

```
switch# show islb merge status
Merge Status: Success
```

Merge conflicts may occur. User intervention is required for the following merge conflicts:

- The iSCSI global authentication or iSCSI initiator idle timeout parameters are not configured the same in the two fabrics.
- The same iSLB initiator is configured differently in the two fabrics.
- An iSLB initiator in one fabric has the same name as an iSCSI initiator in the other fabric.
- Duplicate pWWN/nWWN configuration is detected in the two fabric. For example, a pWWN/nWWN configured for an iSLB initiator on one fabric is configured for an iSCSI initiator or a different iSLB initiator in the other fabric.
- A VSAN configured for an iSLB initiator in one fabric does not exist in the other fabric.

Tip Check the syslog for details on merge conflicts.

User intervention is not required when the same iSLB initiator has a different set of non-conflicting initiator targets. The merged configuration is the union of all the initiator targets.

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- Transparent Target Failover, page 42-61
- Multiple IPS Ports Connected to the Same IP Network, page 42-66
- VRRP-Based High Availability, page 42-67
- Ethernet PortChannel-Based High Availability, page 42-68

Transparent Target Failover

The following high availability configurations are available:

- iSCSI high availability with host running multi-path software
- iSCSI High availability with host not having multi-path software

iSCSI High Availability with Host Running Multi-Path Software

Figure 42-15 shows the physical and logical topology for an iSCSI HA solution for hosts running multi-path software. In this scenario, the host has four iSCSI sessions. There are two iSCSI sessions from each host NIC to the two IPS ports.

Figure 42-15 Host Running Multi-Path Software

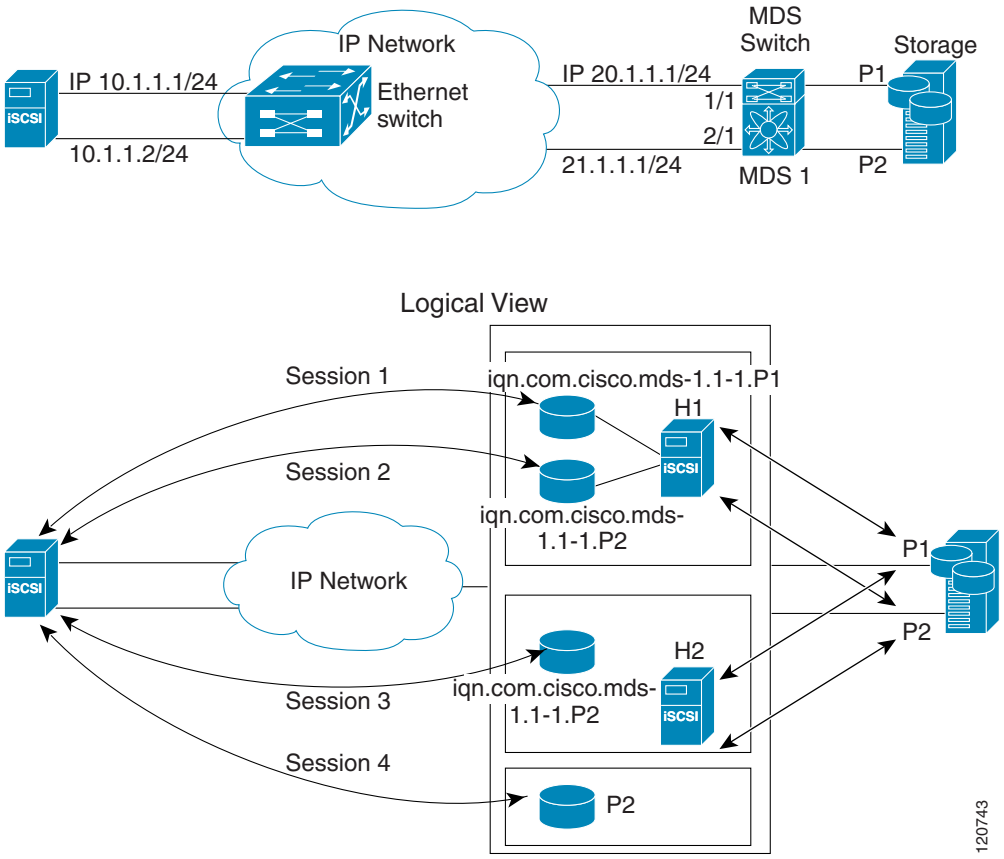

Each IPS ports is exporting the same two Fibre Channel target ports of the storage but as different iSCSI target names if you use dynamic iSCSI targets). So the two IPS ports are exporting a total of four iSCSI target devices. These four iSCSI targets map the same two ports of the Fibre Channel target.

The iSCSI host uses NIC-1 to connect to IPS port 1 and NIC-2 to connect to IPS port 2. Each IPS port exports two iSCSI targets, so the iSCSI host creates four iSCSI sessions.

If the iSCSI host NIC-1 fails (see Figure 42-15 for the physical view), then sessions 1 and 2 fail but we still have sessions 3 and 4.

If the IPS port 1 fails, the iSCSI host cannot connect to the IPS port, and sessions 1 and 2 fail. But sessions 3 and 4 are still available.

If the storage port 1 fails, then the IPS ports will terminate sessions 1 and 3 (put iSCSI virtual target iqn.com.cisco.mds-5.1-2.p1 and iqn-com.cisco.mds-5.1-1.p1 in offline state). But sessions 2 and 4 are still available.

In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load-balancing or failover across the different paths to access the storage.

iSCSI HA with Host Not Having Any Multi-Path Software

The above topology will not work if the host does not have multi-path software because the host has multiple sessions to the same storage. Without multi-path software the host does not have knowledge of the multiple paths to the same storage.

IP storage has two additional features that provide an HA solution in this scenario.

- IPS ports support the VRRP feature (see the "Configuring VRRP for Gigabit Ethernet Interfaces" section on page 44-6) to provide failover for IPS ports.
- IPS has transparent Fibre Channel target failover for iSCSI static virtual targets.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see Figure 42-16).

Figure 42-16 Static Target Importing Through Two Fibre Channel Ports

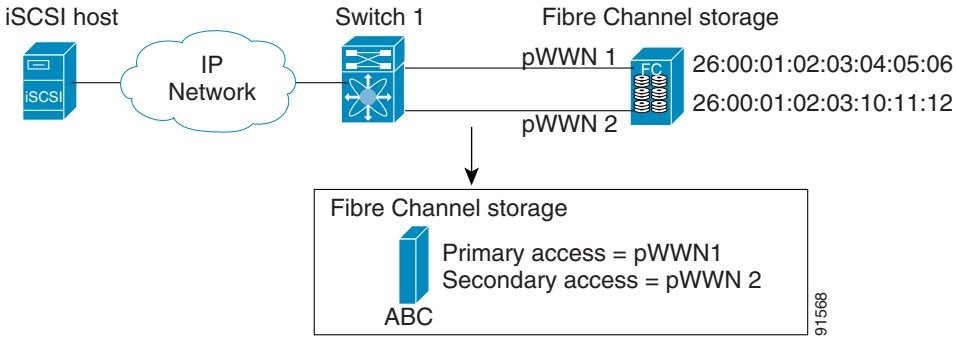

In Figure 42-16, you can create an iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

The failover to a secondary port is done transparently by the IPS port without impacting the iSCSI session from the host. All outstanding I/Os are terminated with a check condition status when the primary port fails. New I/Os received during the failover are not completed and receive a busy status.

Tip If you use LUN mapping, you can define a different secondary Fibre Channel LUN if the LU number is different.

Enable the optional **revert-primary-port** option to direct the IPS port to switch back to the primary port when the primary port is up again. If this option is disabled (default) and the primary port is up again after a switchover, the old sessions will remain with the secondary port and do not switch back to the primary port. However, any new session will use the primary port. This is the only situation when both the primary and secondary ports are used at the same time.

To create a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator`	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.

	Command	Purpose
Step 3	`switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06`	Configures the primary port for this virtual target.
	`switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 secondary-pwwn 26:00:01:02:03:10:11:12`	Configures the primary and secondary ports for this virtual target.
	`switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 fc-lun 0x1 iscsi-lun 0x0 sec-lun 0x3`	Configures the primary port for this virtual target with LUN mapping and different LUN on the secondary Fibre Channel port. **Note** The CLI interprets the LUN identifier value as a hexadecimal value whether or not the **0x** prefix is included.
	`switch(config-iscsi-tgt)# no pwwn 26:00:01:02:03:04:05:06`	Removes the primary port, secondary port, and LUN mapping configuration for this virtual target.
Step 4	`switch(config-iscsi-tgt)# revert-primary-port`	Configures the session failover redundancy for this virtual-target to switch all sessions back to primary port when the primary port comes back up.
Step 5	`switch(config-iscsi-tgt)# no revert-primary-port`	Directs the switch to continue using the secondary port for existing sessions and to use the primary port for new sessions (default).

LUN Trespass for Storage Port Failover

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available to enable the move of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the trespass feature be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see Figure 42-17).

Figure 42-17 Virtual Target with an Active Primary Port

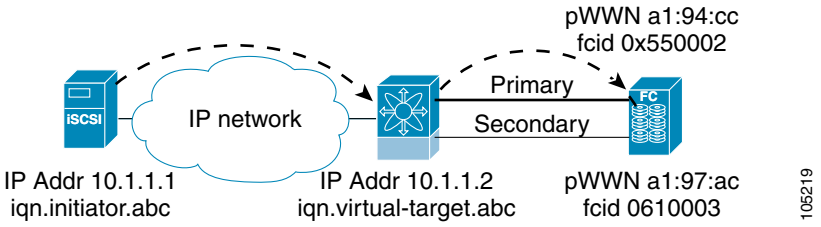

To enable the trespass feature for a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator` `switch(config-iscsi-tgt)#`	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
Step 3	`switch(config-iscsi-tgt)# pwwn 50:00:00:a1:94:cc secondary-pwwn 50:00:00:a1:97:ac`	Maps a virtual target node to a Fibre Channel target and configures a secondary pWWN.
Step 4	`switch(config-iscsi-tgt)# trespass`	Enables the trespass feature.
	`switch(config-iscsi-tgt)# no trespass`	Disables the trespass feature (default).

Use the **show iscsi virtual-target** command to verify.

```
switch# show iscsi virtual-target iqn.1987-02.com.cisco.initiator
target: 1987-02.com.cisco.initiator
    Port WWN 10:20:10:00:56:00:70:50
    Configured node
    all initiator permit is disabled
    trespass support is enabled
```

Multiple IPS Ports Connected to the Same IP Network

Figure 42-18 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 42-18 Multiple Gigabit Ethernet Interfaces in the Same IP Network

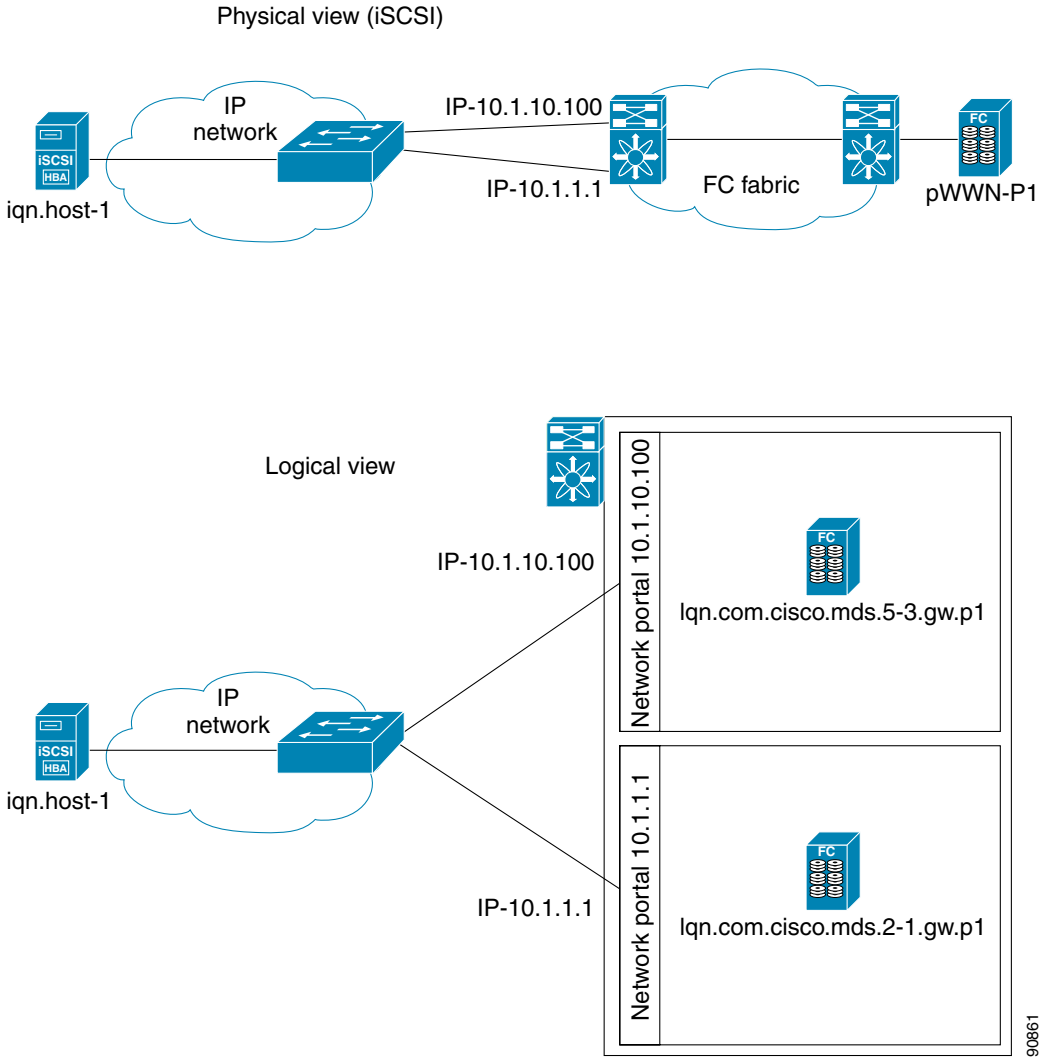

In Figure 42-18, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 42-19 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 42-19 VRRP-Based iSCSI High Availability

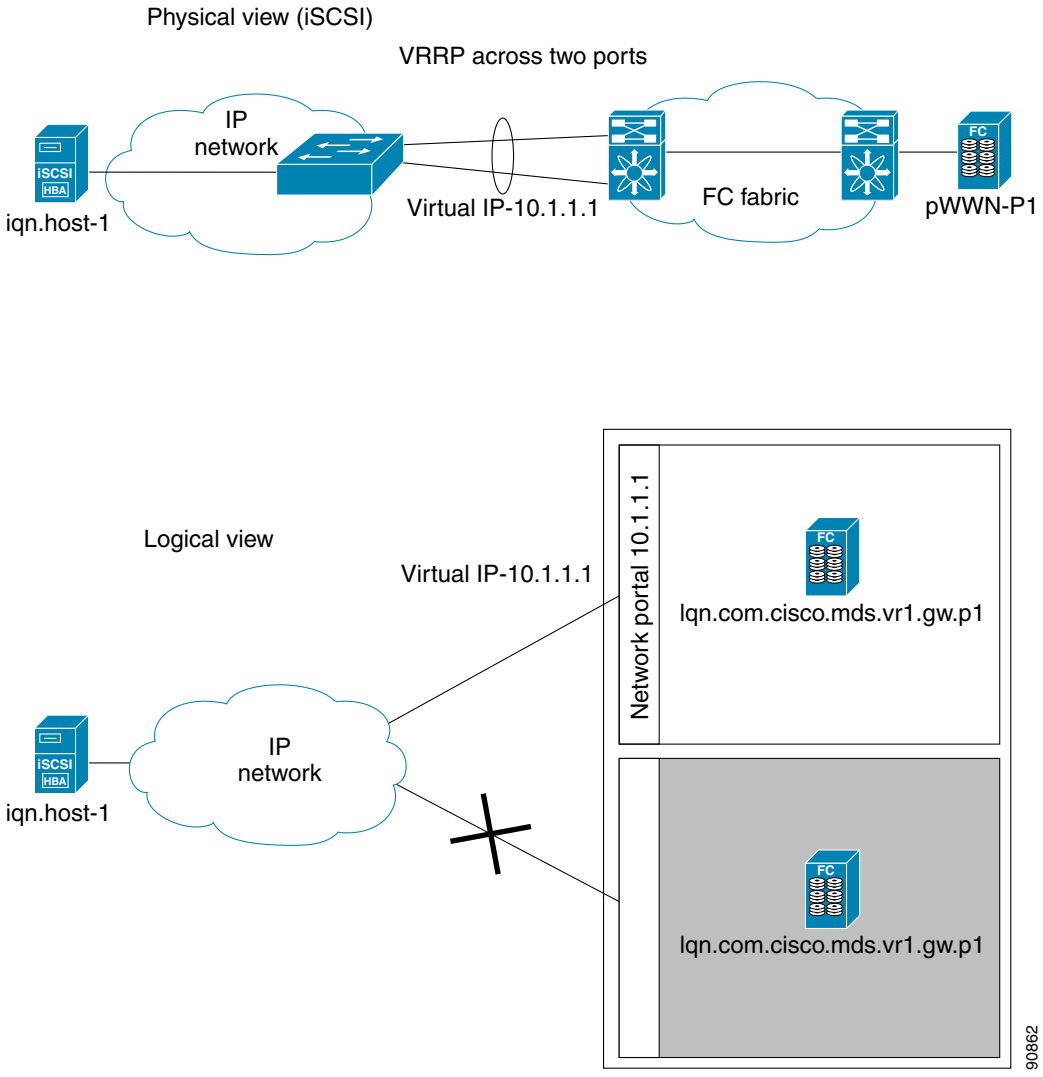

In Figure 42-19, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Ethernet PortChannel-Based High Availability

Note All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that iSCSI link.

Figure 42-20 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 42-20 Ethernet PortChannel-Based iSCSI High Availability

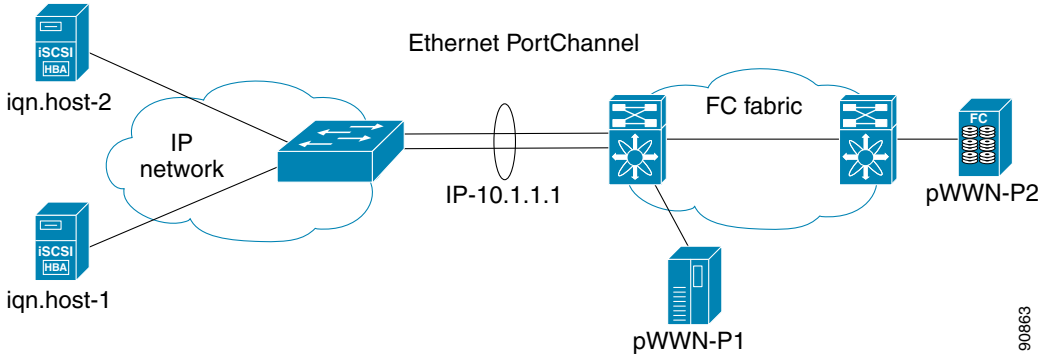

In Figure 42-20, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

Note If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.

iSCSI Authentication Setup Guidelines and Scenarios

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- No Authentication, page 42-69
- CHAP with Local Password Database, page 42-69
- CHAP with External RADIUS Server, page 42-70
- iSCSI Transparent Mode Initiator, page 42-71
- Target Storage Device Requiring LUN Mapping, page 42-76

Note This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before issuing any command.

Caution Changing the authentication of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the "Changing iSCSI Interface Parameters and the Impact on Load Balancing" section on page 42-53.

No Authentication

Set the iSCSI authentication method to **none** to configure a network with no authentication.

```
switch(config)# iscsi authentication none
```

CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

Step 1 Set the AAA authentication to use the local password database for the iSCSI protocol.

```
switch(config)# aaa authentication iscsi default local
```

Step 2 Set the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

Step 3 Configure the user names and passwords for iSCSI users.

```
switch(config)# username iscsi-user password abcd iscsi
```

Note If you do not specify the **iscsi** option, the user name is assumed to be a Cisco MDS switch user instead of an iSCSI user.

Step 4 Verify the global iSCSI authentication setup.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----Verify
  Import FC Target: Disabled
  ...
```

CHAP with External RADIUS Server

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

Step 1 Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server:

```
switch(config)# radius-server key mds-1
```

Step 2 Configure the RADIUS server IP address by performing one of the following:

Configure an IPv4 address.

```
switch(config)# radius-server host 10.1.1.10
```

Configure an IPv6 address.

```
switch(config)# radius-server host 2001:0DB8:800:200C::417A
```

Step 3 Configure the RADIUS server group IP address by performing one of the following:

Configure an IPv4 address.

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 10.1.1.1
```

Configure an IPv6 address.

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 001:0DB8:800:200C::4180

switch(config)# aaa authentication iscsi default group iscsi-radius-group
```

Step 4 Set up the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

Step 5 Verify that the global iSCSI authentication setup is for CHAP.

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP        <---------------- Verify CHAP
  ....
```

Step 6 Verify that the AAA authentication information is for iSCSI.

```
switch# show aaa authentication
         default: local
         console: local
         iscsi: group iscsi-radius-group    <--------- Group name
         dhchap: local

switch# show radius-server groups
total number of groups:2

following RADIUS server groups are configured:
        group radius:
                server: all configured radius servers
        group iscsi-radius-group:
                server: 10.1.1.1 on auth-port 1812, acct-port 1813

switch# show radius-server
Global RADIUS shared secret:mds-1    <-------- Verify secret
....

following RADIUS servers are configured:
        10.1.1.1:                    <----------- Verify the server IPv4 address
                available for authentication on port:1812
                available for accounting on port:1813
```

To configure an iSCSI RADIUS server, follow these steps:

Step 1 Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.

Step 2 Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.

Step 3 Configure the iSCSI users and passwords on the RADIUS server.

iSCSI Transparent Mode Initiator

This scenario assumes the following configuration (see Figure 42-21):

- No LUN mapping or LUN masking or any other access control for hosts on the target device
- No iSCSI login authentication (that is, login authentication set to none)

- The topology is as follows:
  - iSCSI interface 7/1 is configured to identify initiators by IP address.
  - iSCSI interface 7/5 is configured to identify initiators by node name.
  - The iSCSI initiator host 1 with IPv4 address 10.11.1.10 and name iqn.1987-05.com.cisco:01.255891611111 connects to IPS port 7/1 is identified using IPv4 address (host 1 = 10.11.1.10).
  - The iSCSI initiator host 2 with IPv4 address 10.15.1.10 and node name iqn.1987-05.com.cisco:01.25589167f74c connects to IPS port 7/5.

Figure 42-21 iSCSI Scenario 1

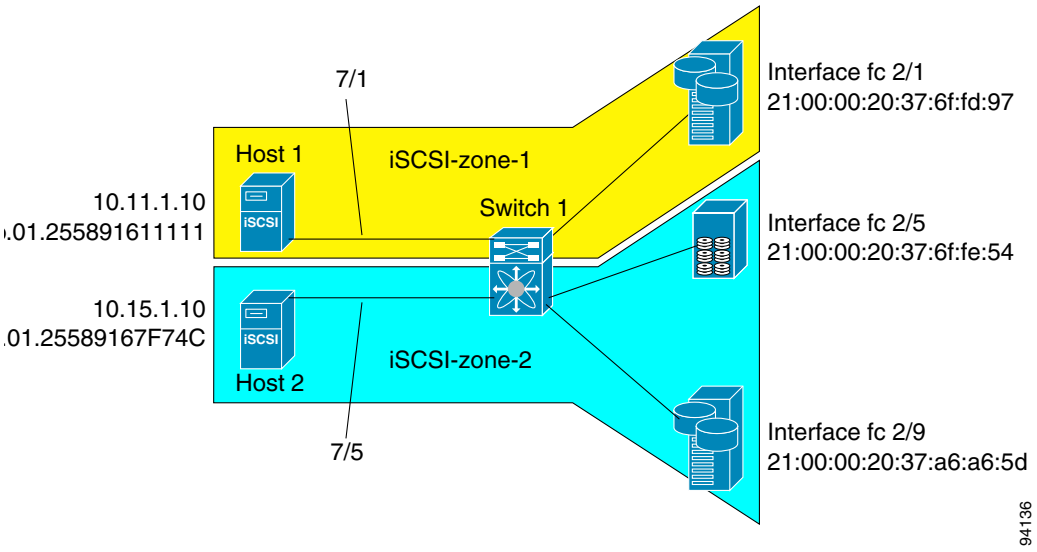

To configure scenario 1 (see Figure 42-21), follow these steps:

Step 1 Configure null authentication for all iSCSI hosts in Cisco MDS switches.

```
switch(config)# iscsi authentication none
```

Step 2 Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

```
switch(config)# iscsi import target fc
```

Step 3 Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

```
switch(config)# interface gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shutdown
```

Note Host 2 is connected to this port.

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address, and enable the interface.

```
switch(config)# interface iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```

Step 5 Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.

```
switch(config)# interface gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shutdown
```

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.

```
switch(config)# interface iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shutdown
```

Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets (see Figure 42-21).

```
switch# show fcns database
VSAN 1:
-----------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
-----------------------------------------------------------------------
0x6d0001    NL    21:00:00:20:37:6f:fd:97 (Seagate)       scsi-fcp:target
0x6d0101    NL    21:00:00:20:37:6f:fe:54 (Seagate)       scsi-fcp:target
0x6d0201    NL    21:00:00:20:37:a6:a6:5d (Seagate)       scsi-fcp:target
Total number of entries = 3
```

Step 8 Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.

Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

```
switch(config)# zone name iscsi-zone-1 vsan 1
switch(config-zone)# member pwwn 21:00:00:20:37:6f:fd:97
switch(config-zone)# member ip-address 10.11.1.10
```

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two Fibre Channel targets in it.

Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

```
switch(config)# zone name iscsi-zone-2 vsan 1
switch(config-zone)# member pwwn 21:00:00:20:37:6f:fe:54
switch(config-zone)# member pwwn 21:00:00:20:37:a6:a6:5d
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c
```

Step 10 Create a zone set and add the two zones as members.

```
switch(config)# zoneset name zoneset-iscsi vsan 1
switch(config-zoneset)# member iscsi-zone-1
switch(config-zoneset)# member iscsi-zone-2
```

Step 11 Activate the zone set.

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

Step 12 Display the active zone set.

Note The iSCSI hosts are not connected so they do not have an FC ID yet.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
  * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97] <----------Target
    symbolic-nodename 10.11.1.10 <---------- iSCSI host (host 1, not online)

  zone name iscsi-zone-2 vsan 1
  * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54] <----------Target
  * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d] <----------Target
    symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c <-iSCSI host (host 2, not online)
```

Step 13 Bring up the iSCSI hosts (host 1 and host 2).

Step 14 Show all the iSCSI sessions (use the **detail** option for detailed information).

```
switch# show iscsi session
    Initiator iqn.1987-05.com.cisco:01.25589167f74c <-----Host 2
    Initiator ip addr (s): 10.15.1.11
  Session #1
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54
```

Note The last part of the auto-created target name is the Fibre Channel target's pWWN.

```
    VSAN 1, ISID 00023d000001, Status active, no reservation

  Session #2
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d
    VSAN 1, ISID 00023d000001, Status active, no reservation

Initiator 10.11.1.10 <-----------------------------------Host 1
  Initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
  Session #1
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97
    VSAN 1, ISID 00023d000001, Status active, no reservation
```

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.25589167f74c <---------
    Initiator ip addr (s): 10.15.1.11
    iSCSI alias name: oasis11.cisco.com
    Node WWN is 20:02:00:0b:fd:44:68:c2 (dynamic)
    Member of vsans: 1
    Number of Virtual n_ports: 1
    Virtual Port WWN is 20:03:00:0b:fd:44:68:c2 (dynamic)
      Interface iSCSI 7/5, Portal group tag: 0x304
      VSAN ID 1, FCID 0x6d0300

iSCSI Node name is 10.11.1.10  <------------------------
    iSCSI Initiator name: iqn.1987 - 05.com.cisco:01.e41695d16b1a
    iSCSI alias name: oasis10.cisco.com
    Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic)
    Member of vsans: 1
    Number of Virtual n_ports: 1
    Virtual Port WWN is 20:05:00:0b:fd:44:68:c2 (dynamic)
      Interface iSCSI 7/1, Portal group tag: 0x300
      VSAN ID 1, FCID 0x6d0301
```

Host 2: Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Host 1: Initiator ID based on IPv4 address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FC IDs are resolved.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
  * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97]
  * fcid 0x6d0301 [symbolic-nodename 10.11.1.10] <-------------

 zone name iscsi-zone-2 vsan 1
  * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54]
  * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d]
  * fcid 0x6d0300 [symbolic-nodename
iqn.1987-05.com.cisco:01.25589167f74c]<-------------------------
```

FC ID resolved for host 1

FC ID for host 2

Step 17 The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

```
switch# show fcns database
VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x6d0001    NL    21:00:00:20:37:6f:fd:97 (Seagate)       scsi-fcp:target
0x6d0101    NL    21:00:00:20:37:6f:fe:54 (Seagate)       scsi-fcp:target
0x6d0201    NL    21:00:00:20:37:a6:a6:5d (Seagate)       scsi-fcp:target
0x6d0300    N     20:03:00:0b:fd:44:68:c2 (Cisco)         scsi-fcp:init isc..w
0x6d0301    N     20:05:00:0b:fd:44:68:c2 (Cisco)         scsi-fcp:init isc..w
```

Step 18 Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

```
switch# show fcns database fcid 0x6d0300 detail vsan 1
------------------------
VSAN:1     FCID:0x6d0300
------------------------
port-wwn (vendor)     :20:03:00:0b:fd:44:68:c2 (Cisco)
node-wwn              :20:02:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr          :10.15.1.11    <----------------------
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw   <------
symbolic-port-name    :

symbolic-node-name
:iqn.1987-05.com.cisco:01.25589167f74c<--------------------
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:91:00:0b:fd:44:68:c0
hard-addr             :0x000000
Total number of entries = 1

switch# show fcns database fcid 0x6d0301 detail vsan 1
------------------------
VSAN:1     FCID:0x6d0301
------------------------
port-wwn (vendor)     :20:05:00:0b:fd:44:68:c2 (Cisco)
node-wwn              :20:04:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr          :10.11.1.10
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw   <------
symbolic-port-name    :

symbolic-node-name    :10.11.1.10  <--------------------
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:81:00:0b:fd:44:68:c0
hard-addr             :0x000000
```

IPv4 address of the iSCSI host

iSCSI gateway node

iSCSI initiator ID is based on the registered node name

iSCSI gateway node

iSCSI initiator ID is based on the IPv4 address registered in symbolic-node-name field

Target Storage Device Requiring LUN Mapping

Sample scenario 2 assumes the following configuration (see Figure 42-22):

- Access control is based on Fibre Channel zoning.
- There is target-based LUN mapping or LUN masking.
- There is no iSCSI authentication (none).
- The iSCSI initiator is assigned to different VSANs.

Figure 42-22 iSCSI Scenario 2

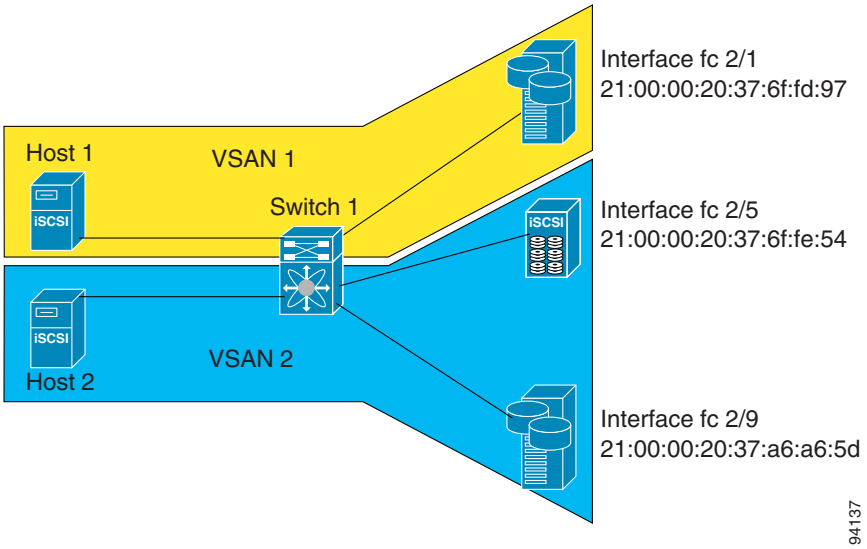

To configure scenario 2 (see Figure 42-22), follow these steps:

Step 1 Configure null authentication for all iSCSI hosts.

```
switch(config)# iscsi authentication none
```

Step 2 Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

```
switch(config)# iscsi import target fc
```

Step 3 Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.

```
switch(config)# interface gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shutdown
```

Step 4 Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address and enable the interface.

```
switch(config)# interface iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shutdown
```

Step 5 Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.

```
switch(config)# interface gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shutdown
```

Step 6 Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.

```
switch(config)# interface iscsi 7/5
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shutdown
```

Step 7 Add static configuration for each iSCSI initiator.

```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a <-----Host 2
switch(config-iscsi-init)# static pWWN system-assign 1
switch(config-iscsi-init)# static nWWN system-assign

switch(config)# iscsi initiator ip address 10.15.1.11 <----------------------------------Host 1
switch(config-iscsi-init)# static pwwn system-assigned 1
switch(config-iscsi-init)# vsan 2
```

Note Host 1 is configured in VSAN 2.

Step 8 View the configured WWNs.

Note The WWNs are assigned by the system. The initiators are members of different VSANs.

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
    Member of vsans: 1
    Node WWN is 20:03:00:0b:fd:44:68:c2
    No. of PWWN: 1
      Port WWN is 20:02:00:0b:fd:44:68:c2

iSCSI Node name is 10.15.1.11
    Member of vsans: 2
    No. of PWWN: 1
      Port WWN is 20:06:00:0b:fd:44:68:c2
```

Step 9 Create a zone with host 1.

```
switch(config)# zone name iscsi-zone-1 vsan 1
```

Step 10 Add three members to the zone named *iscsi-zone-1*.

Note Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN, can be used. In this case, the pWWN is persistent.

- The following command is based on the symbolic node name.

  ```
  switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
  ```

- The following command is based on the persistent pWWN assigned to the initiator. You can obtain the pWWN from the **show iscsi initiator** output.

  ```
  switch(config-zone)# member pwwn 20:02:00:0b:fd:44:68:c2
  ```

Step 11 Create a zone with host 2 and two Fibre Channel targets.

Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

```
switch(config)# zone name iscsi-zone-2 vsan 2
```

Step 12 Activate the zone set in VSAN 2

```
switch(config)# zoneset activate name iscsi-zoneset-v2 vsan 2
Zoneset activation initiated. check zone status
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
  * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
  * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
    pwwn 20:06:00:0b:fd:44:68:c2      <-----------------Host is not online
```

Step 13 Start the iSCSI clients on both hosts and verify that sessions come up.

Step 14 Display the iSCSI sessions to verify the Fibre Channel target and the configured WWNs.

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.e41695d16b1a
    Initiator ip addr (s): 10.11.1.10
  Session #1
    Discovery session, ISID 00023d000001, Status active

  Session #2
    Target
iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97<----
    VSAN 1, ISID 00023d000001, Status active, no reservation
```

To Fibre Channel target

Step 15 Display the iSCSI initiator to verify the configured nWWN and pWWN.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
    Initiator ip addr (s): 10.11.1.10
    iSCSI alias name: oasis10.cisco.com

    Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----------
    Member of vsans: 1
    Number of Virtual n_ports: 1

 Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<----
      Interface iSCSI 7/1, Portal group tag: 0x300
      VSAN ID 1, FCID 0x680102
```

The configured nWWN

The configured pWWN

Step 16 Check the Fibre Channel name server.

```
switch# show fcns database vsan 1
VSAN 1:
--------------------------------------------------------------------
FCID     TYPE PWWN                    (VENDOR)  FC4-TYPE:FEATURE
--------------------------------------------------------------------
0x680001 NL   21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102 N    20:02:00:0b:fd:44:68:c2 (Cisco)   scsi-fcp:init iscw <---
```

iSCSI initiator in name server

Step 17 Verify the details of the iSCSI initiator's FC ID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
------------------------
VSAN:1     FCID:0x680102
------------------------
port-wwn (vendor)     :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn              :20:03:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr          :10.11.1.10
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name    :
symbolic-node-name    :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:81:00:0b:fd:44:68:c0
    iSCSI alias name: oasis10.cisco.com
```

Step 18 Check the Fibre Channel name server.

```
switch# show fcns database vsan 1
VSAN 1:
--------------------------------------------------------------------
FCID     TYPE  PWWN                    (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------
0x680001 NL   21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102 N    20:02:00:0b:fd:44:68:c2 (Cisco)   scsi-fcp:init isc..w <-----
```

iSCSI initiator in name server

Step 19 Verify the details of the iSCSI initiator's FC ID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
------------------------
VSAN:1     FCID:0x680102
------------------------
port-wwn (vendor)     :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn              :20:03:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr          :10.11.1.10
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name    :
symbolic-node-name    :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:81:00:0b:fd:44:68:c0
hard-addr             :0x000000
```

Step 20 Verify that zoning has resolved the FC ID for the iSCSI client.

```
switch# show zoneset active vsan 1
zoneset name iscsi-zoneset-v1 vsan 1
  zone name iscsi-zone-1 vsan 1
  * fcid 0x680001 [pwwn 21:00:00:20:37:6f:fd:97]
  * fcid 0x680102 [pwwn 20:02:00:0b:fd:44:68:c2]
```

Step 21 Verify that the second initiator is connected to the two Fibre Channel targets in VSAN 2.

```
switch# show iscsi session initiator 10.15.1.11
Initiator 10.15.1.11
  Initiator name iqn.1987-05.com.cisco:01.25589167f74c
  Session #1
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <--
    VSAN 2, ISID 00023d000001, Status active, no reservation

  Session #2
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d <--
    VSAN 2, ISID 00023d000001, Status active, no reservation
```

Session to first target

Session to second target

```
switch# show iscsi initiator
iSCSI Node name is 10.15.1.11 <--- Initiator ID is the IP address
    iSCSI Initiator name: iqn.1987-05.com.cisco:01.25589167f74c
    iSCSI alias name: oasis11.cisco.com

    Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <-----------------------
    Member of vsans: 2 <--- vsan membership
    Number of Virtual n_ports: 1

    Virtual Port WWN is 20:06:00:0b:fd:44:68:c2 (configured) <---------
      Interface iSCSI 7/5, Portal group tag: 0x304
      VSAN ID 2, FCID 0x750200
```

Dynamic WWN as static WWN not assigned

Static pWWN for the initiator

```
switch# show fcns database vsan 2
VSAN 2:
--------------------------------------------------------------------------
FCID        TYPE  PWNN                    (VENDOR)  FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x750001    NL    21:00:00:20:37:6f:fe:54 (Seagate) scsi-fcp:target
0x750101    NL    21:00:00:20:37:a6:a6:5d (Seagate) scsi-fcp:target

0x750200    N     20:06:00:0b:fd:44:68:c2 (Cisco)   scsi-fcp:init isc..w <--
Total number of entries = 3
```

iSCSI initiator entry in name server

```
switch# show fcns database fcid 0x750200 detail vsan 2
------------------------
VSAN:2     FCID:0x750200
------------------------
port-wwn (vendor)     :20:06:00:0b:fd:44:68:c2 (Cisco)
node-wwn              :20:04:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr          :10.15.1.11
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name    :
symbolic-node-name    :10.15.1.11
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:91:00:0b:fd:44:68:c0
hard-addr             :0x000000
Total number of entries = 1
```

```
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
  * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
  * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]

 * fcid 0x750200 [pwwn 20:06:00:0b:fd:44:68:c2] <------------------------
```

FC ID resolved for iSCSI initiator

iSNS

Internet Storage Name Service (iSNS) allows your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. To facilitate these functions, the iSNS server and client function as follows:

- The iSNS client registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server.
- The iSNS server provides the following services for the iSNS client:
  - Device registration
  - State change notification
  - Remote domain discovery services

All iSCSI devices (both initiator and target) acting as iSNS clients, can register with an iSNS server. iSCSI initiators can then query the iSNS server for a list of targets. The iSNS server will respond with a list of targets that the querying client can access based on configured access control parameters.

A Cisco MDS 9000 Family switch can act as an iSNS client and register all available iSCSI targets with an external iSNS server. All switches in the Cisco MDS 9000 Family with IPS modules or MPS-14/2 modules installed support iSNS server functionality. This allows external iSNS clients, such as an iSCSI initiator, to register with the switch and discover all available iSCSI targets in the SAN.

This section includes the following topics:

About iSNS Client Functionality

The iSNS client functionality on each IPS interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with an iSNS server. You must specify an iSNS server's IP address by creating an iSNS profile, adding the server's IP address to it, and then assigning (or "tagging") the profile to the interface. An iSNS profile can be tagged to one or more interfaces.

Once a profile is tagged to an interface, the switch opens a TCP connection to the iSNS server IP address (using the well-known iSNS port number 3205) in the profile and registers network entity and portal objects; a unique entity is associated with each IPS interface. The switch then searches the Fibre Channel name server (FCNS) database and switch configuration to find storage nodes to register with the iSNS server.

Statically mapped virtual targets are registered if the associated Fibre Channel pWWN is present in the FCNS database and no access control configuration prevents it. A dynamically mapped target is registered if dynamic target importing is enabled. See the "Presenting Fibre Channel Targets as iSCSI Targets" section on page 42-6 for more details on how iSCSI imports Fibre Channel targets.

A storage node is deregistered from the iSNS server when it becomes unavailable when a configuration changes (such as access control change or dynamic import disabling) or the Fibre Channel storage port goes offline. It is registered again when the node comes back online.

When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server. The iSNS client uses a registration interval value of 15 minutes. If the client fails to refresh the registration during this interval, the server will deregister the entries.

Untagging a profile also causes the network entity and portal to be deregistered from that interface.

Note The iSNS client is not supported on a VRRP interface.

Creating an iSNS Client Profile

To create an iSNS profile, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# isns profile name MyIsns` `switch(config-isns-profile)#`	Creates a profile called MyIsns.
Step 3	`switch(config-isns-profile)# server 10.10.100.211`	Specifies an iSNS server IPv4 address for this profile.
Step 4	`switch(config-isns-profile)# no server 10.10.100.211`	Removes a configured iSNS server from this profile.
Step 5	`switch(config-isns-profile)# server 2003::11`	Specifies an iSNS server IPv6 address for this profile.
Step 6	`switch(config-isns-profile)# no server 10.20.100.211`	Removes a configured iSNS server from this profile.

To remove an iSNS profile, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no isns profile name OldIsns`	Removes a configured iSNS profile called OldIsns.

To tag a profile to an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 4/1` `switch(config-if)#`	Configures the specified Gigabit Ethernet interface.
Step 3	`switch(config-if)# isns MyIsns`	Tags a profile to an interface.

To untag a profile from an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 5/1` `switch(config-if)#`	Configures the specified Gigabit Ethernet interface.
Step 3	`switch(config-if)# no isns OldIsns`	Untags a profile from an interface.

Use the **isns reregister** command in EXEC mode to re-register associated iSNS objects with the iSNS server.

```
switch# isns reregister gigabitethernet 1/4
switch# isns reregister port-channel 1
```

Verifying iSNS Client Configuration

Use the **show isns profile** command to view configured iSNS profiles. Profile ABC has two portals registered with the iSNS server. Each portal corresponds to a particular interface. Profile XYZ has a specified iSNS server, but does not have any tagged interfaces configured (see Example 42-19 and Example 42-20).

Example 42-19 Displays Information for Configured iSNS Profiles

```
switch# show isns profile
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204

iSNS profile name XYZ
iSNS Server 10.10.100.211
```

Example 42-20 Displays a Specified iSNS Profile

```
switch# show isns profile ABC
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204
```

Use the **show isns profile counters** command to view all configured profiles with the iSNS PDU statistics for each tagged interface (see Example 42-21 and Example 42-22).

Example 42-21 Displays Configured Profiles with iSNS Statistics

```
switch# show isns profile counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
    Reg pdus 37,  Dereg pdus 17
  Output 54 pdus (registration/deregistration pdus only)
    Reg pdus  37,  Dereg pdus 17
iSNS Server  10.10.100.204

iSNS profile name XYZ
tagged interface  port-channel 2
iSNS statistics
  Input 30 pdus (registration/deregistration pdus only)
    Reg pdus 29,  Dereg pdus 1
  Output 30 pdus (registration/deregistration pdus only)
    Reg pdus  29,  Dereg pdus 1
iSNS Server  10.1.4.218
```

Example 42-22 Displays iSNS Statistics for a Specified Profile

```
switch# show isns profile ABC counters
iSNS profile name ABC
tagged interface  port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
    Reg pdus 37,  Dereg pdus 17
  Output 54 pdus (registration/deregistration pdus only)
    Reg pdus  37,  Dereg pdus 17
iSNS Server  10.10.100.204
```

Use the **show isns** command to view all objects registered on the iSNS server and specified in the given profile (see Example 42-23).

Example 42-23 Displays iSNS Queries

```
switch# show isns query ABC gigabitethernet 2/3
iSNS server: 10.10.100.204
Init: iqn.1991-05.com.w2k
  Alias: <MS SW iSCSI Initiator>
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03.210000203762fa34
  nWWN: 200000203762fa34
```

Use the **show interface** command to view the iSNS profile to which an interface is tagged (see Example 42-24).

Example 42-24 Displays Tagged iSNS Interfaces

```
switch# show interface gigabitethernet 2/3
    GigabitEthernet2/3 is up
    Hardware is GigabitEthernet, address is 0005.3000.ae94
    Internet address is 10.10.100.201/24
    MTU 1500  bytes
```

```
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
iSNS profile ABC
^^^^^^^^^^^^^^^^
5 minutes input rate 112 bits/sec, 14 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1935 packets input, 132567 bytes
  4 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
1 packets output, 42 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors
```

About iSNS Server Functionality

When enabled, the iSNS server on the Cisco 9000 Family MDS switch tracks all registered iSCSI devices. As a result, iSNS clients can locate other iSNS clients by querying the iSNS server. The iSNS server also provides the following functionalities:

- Allows iSNS clients to register, deregister, and query other iSNS clients registered with the iSNS server.
- Provides centralized management for enforcing access control to provide or deny access to targets from specific initiators.
- Provides a notification mechanism for registered iSNS clients to receive change notifications on the status change of other iSNS clients.
- Provides a single access control configuration for both Fibre Channel and iSCSI devices.
- Discovers iSCSI targets that do not have direct IP connectivity to the iSCSI initiators.

Example Scenario

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. Figure 42-23 provides an example of this scenario.

Figure 42-23 Using iSNS Servers in the Cisco MDS Environment

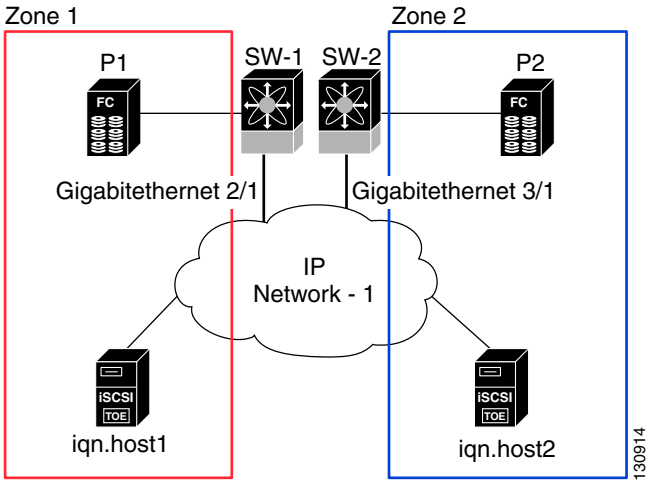

In Figure 42-23, iqn.host1 and iqn.host2 are iSCSI initiators. P1 and P2 are Fibre Channel targets. The two initiators are in different zones: Zone 1 consists of iqn.host1 and target P1, and Zone 2 consists of iqn.host2 and target P2. iSNS server functionality is enabled on both switches, SW-1 and SW-2. The registration process proceeds as follows:

1. Initiator iqn.host1 registers with SW-1, port Gigabitethernet2/1.
2. Initiator iqn.host2 registers with SW-2, port Gigabitethernet3/1.
3. Initiator iqn.host1 issues an iSNS query to SW-1 to determine all accessible targets.
4. The iSNS server in turn queries the Fibre Channel name server (FCNS) to obtain a list of devices that are accessible (that is, in the same zone) by the query originator. This query yields only P1.
5. The iSNS server then queries its own database to convert the Fibre Channel devices to the corresponding iSCSI targets. This is based on the iSCSI configuration, such as virtual-target and its access control setting or whether the dynamic Fibre Channel target import feature is enabled or disabled.
6. The iSNS server sends a response back to the query initiator. This response contains a list all iSCSI portals known to the iSNS server. This means iqn.host1 can choose to log in to target P1 through either SW-1 (at Gigabitethernet 2/1) or SW-2 (at Gigabitethernet 3/1).
7. If the initiator chooses to log in to SW-1 and later that port becomes inaccessible (for example, Gigabitethernet 2/1 goes down), the initiator has the choice to move to connect to target P1 through port Gigabitethernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

Configuring iSNS Servers

This section describe how to configure an iSNS server on a Cisco MDS 9000 Family switch.

This section includes the following topics:

Enabling the iSNS Server

Before the iSNS server feature can be enabled, iSCSI must be enabled (see the "Enabling iSCSI" section on page 42-5). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

To enable the iSNS server, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# isns-server enable`	Enables the iSNS server.
	`switch(config)# no isns-server enable`	Disables (default) the iSNS server.

Note If you are using VRRP IPv4 addresses for discovering targets from iSNS clients, ensure that the IP address is created using the **secondary** option (see the "Adding Virtual Router IP Addresses" section on page 43-20).

iSNS Configuration Distribution

You can use the CFS infrastructure to distribute the iSCSI initiator configuration to iSNS servers across the fabric. This allows the iSNS server running on any switch to provide a querying iSNS client a list of iSCSI devices available anywhere on the fabric. For information on CFS, see Chapter 6, "Using the CFS Infrastructure."

To enable iSNS configuration distribution using, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# isns distribute`	Uses the CFS infrastructure to distribute the iSCSI virtual target configuration to all switches in the fabric.
	`switch(config)# no isns distribute`	Stops (default) the distribution of iSCSI virtual target configuration to all switches in the fabric.

Configuring the ESI Retry Count

The iSNS client registers information with its configured iSNS server using an iSNS profile. At registration, the client can indicate an entity status inquiry (ESI) interval of 60 seconds or more. If the client registers with an ESI interval set to zero (0), then the server does not monitor the client using ESI. In such cases, the client's registrations remain valid until explicitly deregistered or the iSNS server feature is disabled.

The ESI retry count is the number of times the iSNS server queries iSNS clients for their entity status. The default ESI retry count is 3. The client sends the server a response to indicate that it is still alive. If the client fails to respond after the configured number of retries, the client is deregistered from the server.

To configure the ESI retry count for an iSNS server, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# isns esi retries 6`	Configures the ESI to retry contacting the client up to 6 times. The range is 1 to 10.
	`switch(config)# no isns esi retries 6`	Reverts to the default value of 3 retries.

Configuring the Registration Period

The iSNS client specifies the registration period with the iSNS Server. The iSNS Server keeps the registration active until the end of this period. If there are no commands from the iSNS client during this period, then the iSNS server removes the client registration from its database.

If the iSNS client does not specify a registration period, the iSNS server assumes a default value of 0, which keeps the registration active indefinitely. You can also manually configure the registration period on the MDS iSNS Server.

To configure the registration period on an iSNS Server, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# isns registration period 300`	Configures the registration to be active for 300 seconds. The permissible registration period is between 0 - 65536 seconds.
	`switch(config)# no isns registration period`	Reverts to the client registered timeout value, or the default value of 0.

iSNS Client Registration and Deregistration

An iSNS client cannot query the iSNS server until it has registered. You can use the **show isns database** command to display all registered iSNS clients and their associated configuration.

iSNS client deregistration can occur either explicitly or when the iSNS server detects that it can no longer reach the client (through ESI monitoring).

iSNS client registration and deregistration result in status change notifications (SCNs) being generated to all interested iSNS clients.

Target Discovery

iSCSI initiators discover targets by issuing queries to the iSNS server. The server supports *DevGetNext* requests to search the list of targets and *DevAttrQuery* to determine target and portal details, such as the IP address or port number to which to connect.

On receiving a query request from the iSCSI client, the iSNS server queries the Fibre Channel Name Server (FCNS) to obtain a list of Fibre Channel targets that are accessible by the querying initiator. The result of this query depends on zoning configuration currently active and current configuration(s) of the

initiator. The iSNS server will subsequently use the iSCSI target configuration(s) (virtual target and dynamic import configuration) to translate the Fibre Channel target to an equivalent iSCSI target. At this stage it also applies any access control configured for the virtual target. A response message with the target details is then sent back to the query initiator.

The iSNS server sends a consolidated response containing all possible targets and portals to the querying initiator. For example, if a Fibre Channel target is exported as different iSCSI targets on different IPS interfaces, the iSNS server will respond with a list of all possible iSCSI targets and portals.

In order to keep the list of targets updated, the iSNS server sends state change notifications (SCN) to the client whenever an iSCSI target becomes reachable or unreachable. The client is then expected to rediscover its list of accessible targets by initiating another iSNS query. Reachability of iSCSI targets changes when any one of the following occurs:

- Target goes up or down.
- Dynamic import of FC target configuration changes.
- Zone set changes.
- Default zone access control changes.
- IPS interface state changes.
- Initiator configuration change makes the target accessible or inaccessible.

Verifying the iSNS Server Configuration

Use the **show isns config** command to view the ESI interval and the summary information about the iSNS database contents (see Example 42-25).

Example 42-25 Displays the iSNS Server Configuration of ESI Interval and Database Contents

```
switch# show isns config
Server Name: switch1(Cisco Systems) Up since: Fri Jul 30 04:08:16 2004
    Index: 1    Version: 1     TCP Port: 3205
    fabric distribute (remote sync): ON
    ESI
       Non Response Threshold: 5 Interval(seconds): 60
   Database contents
       Number of Entities: 2
       Number of Portals: 3
       Number of ISCSI devices: 4
       Number of Portal Groups: 0
```

Use the **show isns database** command to view detailed information about the contents of the iSNS database (see Example 42-26 through Example 42-29). This command displays the full iSNS database giving all the entities, nodes, and portals registered in the database. This command without options only displays explicitly registered objects. The asterisk next to the VSAN ID indicates that the iSCSI node is in the default zone for that VSAN.

Example 42-26 Displays Explicitly Registered Objects

```
switch# show isns database
Entity Id: dp-204
    Index: 2             Last accessed: Fri Jul 30 04:08:46 2004

iSCSI Node Name: iqn.1991-05.comdp-2041
    Entity Index: 2
    Node Type: Initiator(2)      Node Index: 0x1
    SCN Bitmap: OBJ_UPDATED|OBJ ADDED|OBJ REMOVED|TARGET&SELF
    Node Alias: <MS SW iSCSI Initiator>

    VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2          TCP Port: 4179
    Entity Index: 2    Portal Index: 1
    ESI Interval: 0     ESI Port: 4180     SCN Port: 4180
```

Example 42-27 displays information about both virtual and registered iSCSI initiators/targets.

Example 42-27 Displays the Full Database with Both Registered and Configured Nodes and Portals

```
switch# show isns database full
Entity Id: isns.entity.mds9000
    Index: 1             Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000001
    WWN(s):
        22:00:00:20:37:39:dc:45
    VSANS:
iSCSI Node Name: iqn.isns-first-virtual-target
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000002

    VSANS:
iSCSI Node Name: iqn.com.cisco.disk2
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000003
    WWN(s):
        22:00:00:20:37:39:dc:45

    VSANS:
Portal IP Address: 192.168.100.5          TCP Port: 3205
    Entity Index: 1    Portal Index: 3

Portal IP Address: 192.168.100.6          TCP Port: 3205
    Entity Index: 1    Portal Index: 5

Entity Id: dp-204
    Index: 2             Last accessed: Fri Jul 30 04:08:46 2004

iSCSI Node Name: iqn.1991-05.com.microsoft:dp-2041
    Entity Index: 2
    Node Type: Initiator(2)      Node Index: 0x1
```

```
    SCN Bitmap: OBJ_UPDATED|OBJ ADDED|OBJ REMOVED|TARGET&SELF
    Node Alias: <MS SW iSCSI Initiator>

    VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2          TCP Port: 4179
    Entity Index: 2    Portal Index: 1
    ESI Interval: 0     ESI Port: 4180     SCN Port: 4180
```

Example 42-28 displays the virtual targets entries on the current switch.

Note The **local option** is only available for virtual targets.

Example 42-28 Displays the Virtual Target Information in the Local Switch

```
switch# show isns database virtual-targets local
Entity Id: isns.entity.mds9000
    Index: 1             Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000001
    WWN(s):
        22:00:00:20:37:39:dc:45

    VSANS:
iSCSI Node Name: iqn.isns-first-virtual-target
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000002

    VSANS:
iSCSI Node Name: iqn.com.cisco.disk2
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000003
    WWN(s):
        22:00:00:20:37:39:dc:45

    VSANS:
Portal IP Address: 192.168.100.5          TCP Port: 3205
    Entity Index: 1    Portal Index: 3

Portal IP Address: 192.168.100.6          TCP Port: 3205
    Entity Index: 1    Portal Index: 5
```

Example 42-29 provides the virtual target information for a specific remote switch. The remote switch is specified using the switch ID (the WWN of the switch).

Example 42-29 Displays Virtual Target for a Specified Switch

```
switch# show isns database virtual-targets switch 20:00:00:0d:ec:01:04:40
Entity Id: isns.entity.mds9000
    Index: 1             Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000001
    WWN(s):
        22:00:00:20:37:39:dc:45

    VSANS:
iSCSI Node Name: iqn.isns-first-virtual-target
```

```
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000002

    VSANS:
iSCSI Node Name: iqn.com.cisco.disk2
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000003
    WWN(s):
        22:00:00:20:37:39:dc:45

    VSANS:
Portal IP Address: 192.168.100.5          TCP Port: 3205
    Entity Index: 1    Portal Index: 3

Portal IP Address: 192.168.100.6          TCP Port: 3205
    Entity Index: 1    Portal Index: 5
```

Use the **show isns node** command to display attributes of nodes registered with the iSNS server (see Example 42-30 through Example 42-32). If you do not specify any options, the server displays the name and node type attribute in a compact format; one per line.

Example 42-30 Displays Explicitly Registered Objects

```
switch# show isns node all
--------------------------------------------------------------------------------
iSCSI Node Name                                                          Type
--------------------------------------------------------------------------------
iqn.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8                   Target
...
iqn.com.cisco.disk1                                                      Target
iqn.com.cisco.ipdisk                                                     Target
iqn.isns-first-virtual-target                                            Target
iqn.1991-05.cw22                                                         Target
iqn.1991-05.cw53                                                         Target
```

Example 42-31 Displays the Specified Node

```
switch# show isns node name iqn.com.cisco.disk1
iSCSI Node Name: iqn.com.cisco.disk1
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000001
    WWN(s):
        22:00:00:20:37:39:dc:45
    VSANS:  1
```

Example 42-32 Displays the Attribute Details for All Nodes

```
switch# show isns node all detail
iSCSI Node Name: iqn.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8f
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x3000003
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40
        WWN(s):
            22:00:00:20:37:5a:6c:8f
        VSANS:  1
...
```

```
iSCSI Node Name: iqn.com.cisco.disk1
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000001
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40
        WWN(s):
            22:00:00:20:37:39:dc:45
        VSANS:  1

iSCSI Node Name: iqn.com.cisco.ipdisk
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000002
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40
        WWN(s):
            22:00:00:20:37:5a:70:1a
        VSANS:  1

iSCSI Node Name: iqn.isns-first-virtual-target
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000003
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: iqn.parna.121212
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000004
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: iqn.parna.121213
    Entity Index: 1
    Node Type: Target(1)      Node Index: 0x80000005
    Configured Switch WWN: 20:00:00:0d:ec:01:04:40
```

Use the **show isns portal** command to display the attributes of a portal along with its accessible nodes (see Example 42-33 through Example 42-37). You can specify portals by using the switch WWN-interface combination or the IP address-port number combination.

Example 42-33 Displays the Attribute Information for All Portals

```
switch# show isns portal all
--------------------------------------------------------------------------------
IPAddress      TCP Port     Index               SCN Port            ESI  port
--------------------------------------------------------------------------------
192.168.100.5    3205         3                   -                   -
192.168.100.6    3205         5                   -                   -
```

Example 42-34 Displays Detailed Attribute Information for All Portals

```
switch# show isns portal all detail
Portal IP Address: 192.168.100.5          TCP Port: 3205
    Entity Index: 1    Portal Index: 3

Portal IP Address: 192.168.100.6          TCP Port: 3205
    Entity Index: 1    Portal Index: 5
```

Example 42-35 Displays Virtual Portals

```
switch# show isns portal virtual
--------------------------------------------------------------------------------
IPAddress      TCP Port     Index               SCN Port            ESI  port
--------------------------------------------------------------------------------
192.168.100.5    3205         3                   -                   -
192.168.100.6    3205         5                   -                   -
```

Example 42-36 Displays Virtual Portals for the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40
--------------------------------------------------------------------------------
IPAddress      TCP Port     Index               SCN Port            ESI  port
--------------------------------------------------------------------------------
192.168.100.5    3205         3                   -                   -
192.168.100.6    3205         5                   -                   -
```

Example 42-37 Displays Detailed Information for the Virtual Portals in the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40 detail
Portal IP Address: 192.168.100.5          TCP Port: 3205
    Entity Index: 1    Portal Index: 3
    Switch WWN: 20:00:00:0d:ec:01:04:40
    Interface: GigabitEthernet2/3

Portal IP Address: 192.168.100.6          TCP Port: 3205
    Entity Index: 1    Portal Index: 5
    Switch WWN: 20:00:00:0d:ec:01:04:40
    Interface: GigabitEthernet2/5
```

Use the **show isns entity** command to display the attributes of an entity along with the list of portals and nodes in that entity (see Example 42-38 through Example 42-42). If you do not specify any option, this command displays the entity ID and number of nodes or portals associated with the entity in a compact format; one per line.

Example 42-38 Displays All Registered Entries

```
switch1# show isns entity
--------------------------------------------------------------------------------
Entity ID                                          Last Accessed
--------------------------------------------------------------------------------
dp-204                                       Tue Sep  7 23:15:42 2004
```

Example 42-39 Displays All Entities in the Database

```
switch# show isns entity all

--------------------------------------------------------------------------------
Entity ID                                                 Last Accessed
--------------------------------------------------------------------------------
isns.entity.mds9000                                 Tue Sep  7 21:33:23 2004

dp-204                                              Tue Sep  7 23:15:42 2004
```

Example 42-40 Displays the Entity with the Specified ID

```
switch1# show isns entity id dp-204
Entity Id: dp-204
    Index: 2             Last accessed: Tue Sep  7 23:15:42 2004
```

Example 42-41 Displays Detailed Information for All Entities in the Database

```
switch1# show isns entity all detail
Entity Id: isns.entity.mds9000
    Index: 1             Last accessed: Tue Sep  7 21:33:23 2004

Entity Id: dp-204
    Index: 2             Last accessed: Tue Sep  7 23:16:34 2004
```

Example 42-42 Displays Virtual Entities

```
switch# show isns entity virtual
Entity Id: isns.entity.mds9000
    Index: 1             Last accessed: Thu Aug  5 00:58:50 2004

Entity Id: dp-204
    Index: 2             Last accessed: Thu Aug  5 01:00:23 2004
```

Use the **show iscsi global config** command to display information about import targets (see Example 42-43 and Example 42-44).

Example 42-43 Displays the Import Target Settings for the Specified Switch

```
switch# show isns iscsi global config switch 20:00:00:05:ec:01:04:00
iSCSI Global configuration:
  Switch: 20:00:00:05:ec:01:04:00 iSCSI Auto Import: Enabled
```

Example 42-44 Displays the Import Target Settings for All Switches

```
switch# show isns iscsi global config all
iSCSI Global configuration:
  Switch: 20:00:44:0d:ec:01:02:40 iSCSI Auto Import: Enabled
```

Use the **show cfs peers** command to display CFS peers switch information about the iSNS application (see Example 42-45).

Example 42-45 Displays the CFS Peer Switch Information for the iSNS Application

```
switch# show cfs peers name isns

Scope      : Physical
--------------------------------------------------
 Switch WWN               IP Address
--------------------------------------------------
 20:00:00:00:ec:01:00:40  10.10.100.11   [Local]

Total number of entries = 1
```

iSNS Cloud Discovery

You can configure iSNS cloud discovery to automate the process of discovering iSNS servers in the IP network.

This section includes the following topics:

About Cloud Discovery

Note iSNS Cloud Discovery is not supported on the Cisco Fabric Switch for IBM BladeCenter and Cisco Fabric Switch for HP c-Class BladeSystem.

When an iSNS server receives a query request, it responds with a list of available targets and the portals through which the initiator can reach the target. The IP network configuration outside the MDS switch may result in only a subset of Gigabit Ethernet interfaces being reachable from the initiator. To ensure that the set of portals returned to the initiator is reachable, the iSNS server needs to know the set of Gigabit Ethernet interfaces that are reachable from a given initiator.

The iSNS cloud discovery feature provides information to the iSNS server on the various interfaces reachable from an initiator by partitioning the interfaces on a switch into disjointed IP clouds. This discovery is achieved by sending messages to all other known IPS ports that are currently up and, depending on the response (or the lack of it), determines if the remote IPS port is in the same IP network or in a different IP network.

Cloud discovery is initiated when the following events occur:

- Manual requests from the CLI initiate cloud discovery from the CLI. This action causes the destruction of existing memberships and makes new ones.
- Auto-discovery of the interface results in an interface being assigned to its correct cloud. All other cloud members are not affected. The membership of each cloud is built incrementally and is initiated by the following events:
  - A Gigabit Ethernet interface comes up. This can be a local or remote Gigabit Ethernet interface.
  - The IP address of a Gigabit Ethernet interface changes.
  - The VRRP configuration on a port changes.

The iSNS server distributes cloud and membership information across all the switches using CFS. Therefore, the cloud membership view is the same on all the switches in the fabric.

Note For CFS distribution to operate correctly for iSNS cloud discovery, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later.

Configuring iSNS Cloud Discovery

This section describes how to configure iSNS cloud discovery and includes the following topics:

Enabling iSNS Cloud Discovery

To enable iSNS cloud discovery, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cloud-discovery enable`	Enables iSNS cloud discovery.
	`switch(config)# no cloud-discovery enable`	Disables (default) iSNS cloud discovery.

Initiating On-Demand iSNS Cloud Discovery

To initiate on-demand iSNS cloud discovery, use the **cloud discover** command in EXEC mode.

The following example shows how to initiate on-demand cloud discovery for the entire fabric:

```
switch# cloud discover
```

Configuring Automatic iSNS Cloud Discovery

To configure automatic iSNS cloud discovery, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cloud discovery auto`	Enables (default) automatic iSNS cloud discovery.
	`switch(config)# no cloud discovery auto`	Disables automatic iSNS cloud discovery.

Verifying Automatic iSNS Cloud Discovery Configuration

To verify the automatic iSNS cloud discovery configuration, use the **show cloud discovery config** command.

```
switch# show cloud discovery config
Auto discovery: Enabled
```

Configuring iSNS Cloud Discovery Distribution

To configure iSNS cloud discovery distribution using CFS, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cloud discovery fabric distribute`	Enables (default) iSNS cloud discovery fabric distribution.
	`switch(config)# no cloud discovery fabric distribute`	Disables iSNS cloud discovery fabric distribution.

Configuring iSNS Cloud Discovery Message Types

You can configure iSNS cloud discovery the type of message to use. By default, iSNS cloud discovery uses ICMP.

To configure iSNS cloud discovery message types, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# cloud discovery message icmp`	Enables (default) iSNS cloud discovery using ICMP messages. **Note** Only ICMP messages are supported.

Verifying Cloud Discovery Status

Use the **show cloud discovery status** command to verify the status of the cloud discovery operation.

```
switch# show cloud discovery status
Discovery status: Succeeded
```

Verifying Cloud Discovery Membership

Use the **show cloud membership all** command to verify the cloud membership for the switch.

```
switch# show cloud membership all
Cloud 2
    GigabitEthernet1/5[20:00:00:0d:ec:02:c6:c0] IP Addr 10.10.10.5
    GigabitEthernet1/6[20:00:00:0d:ec:02:c6:c0] IP Addr 10.10.10.6
  #members=2
```

Use the **show cloud membership unresolved** command to verify the unresolved membership on the switch.

```
switch# show cloud membership unresolved
Undiscovered Cloud
  No members
```

Displaying Cloud Discovery Statistics

Use the **show cloud discovery statistics** command to display the statistics for the cloud discovery operation.

```
switch# show cloud discovery statistics
Global statistics
    Number of Auto Discovery                       = 1
    Number of Manual Discovery                     = 0
    Number of cloud discovery (ping) messages sent = 1
    Number of cloud discovery (ping) success       = 1
```

Default Settings

Table 42-2 lists the default settings for iSCSI parameters.

Table 42-2 Default iSCSI Parameters

Parameters	Default
Number of TCP connections	One per iSCSI session.
minimum-retransmit-time	300 msec.
keepalive-timeout	60 seconds.
max-retransmissions	4 retransmissions.
PMTU discovery	Enabled.
pmtu-enable reset-timeout	3600 sec.

Table 42-2 Default iSCSI Parameters (continued)

Parameters	Default
SACK	Enabled.
max-bandwidth	1Gbps
min-available-bandwidth	70 Mbps.
round-trip-time	1 msec.
Buffer size	4096 KB.
Control TCP and data connection	No packets are transmitted.
TCP congestion window monitoring	Enabled.
Burst size	50 KB.
Jitter	500 microseconds.
TCP connection mode	Active mode is enabled.
Fibre Channel targets to iSCSI	Not imported.
Advertising iSCSI target	Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.
iSCSI hosts mapping to virtual Fibre Channel hosts	Dynamic mapping.
Dynamic iSCSI initiators	Members of the VSAN 1.
Identifying initiators	iSCSI node names.
Advertising static virtual targets	No initiators are allowed to access a virtual target (unless explicitly configured).
iSCSI login authentication	CHAP or none authentication mechanism.
revert-primary-port	Disabled.
Header and data digest	Enabled automatically when iSCSI initiators send requests. This feature cannot be configured and is not available in store-and-forward mode.
iSNS registration interval	60 sec (not configurable).
iSNS registration interval retries	3.
Fabric distribution	Disabled.

Table 42-3 lists the default settings for iSLB parameters.

Table 42-3 Default iSLB Parameters

Parameters	Default
Fabric distribution	Disabled.
Load balancing metric	1000.

CHAPTER 43

Configuring IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMSs):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding or in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

Note For information about configuring IPv6, see Chapter 46, “Configuring IPv6 for Gigabit Ethernet Interfaces.”

This chapter includes the following sections:

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running the IP protocol over an FC interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric (see Figure 43-1).

Figure 43-1 Management Access to Switches

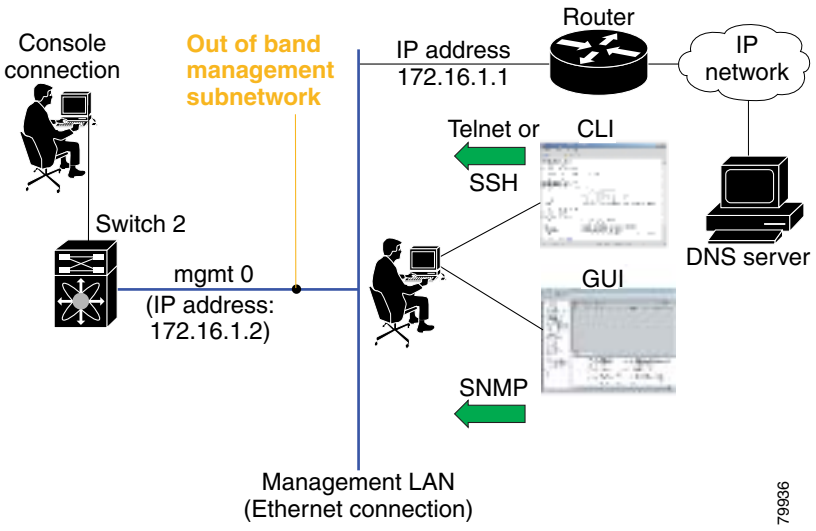

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure IP version 4 (IPv4) parameters (IP address, subnet mask) or an IP version 6 (IPv6) address and prefix length so that the switch is reachable. For information on configuring IPv6 addresses, see Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces."

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.

Note The port on the Ethernet switch to which the MDS management interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should disabled. This helps avoid the delay in the MDS management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in IOS is or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

Note Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

To configure the mgmt0 Ethernet interface for IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	`switch(config-if)# ip address 10.1.1.1 255.255.255.0`	Enters the IPv4 address (10.1.1.1) and IPv4 subnet mask (255.255.255.0) for the management interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

To configure the mgmt0 Ethernet interface for IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	`switch(config-if)# ipv6 address 2001:0db8:800:200c::417a/64`	Enters the IPv6 address (2001:0DB8:800:200C::417A) and IPv6 prefix length (/64) for the management interface and enables IPv6 processing on the interface.
	`switch(config-if)# ipv6 enable`	Automatically configures a link-local IPv6 address on the interface and enables IPv6 processing on the interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Default Gateway

You can configure a default gateway IPv4 address on your Cisco MDS 9000 Family switch.

This section includes the following topics:

About the Default Gateway

The default gateway IPv4 address should be configured along with the IPv4 static routing commands (IP default network, destination prefix, and destination mask, and next hop address).

Tip If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

See the "Initial Setup Routine" section on page 5-2 for more information on configuring the IP addresses for all entries in the switch.

Use the **ip default-gateway** command to configure the IP address for a switch's default gateway and the **show ip route** command to verify that the IPv4 address for the default gateway is configured.

Configuring the Default Gateway

To configure the default gateway, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ip default-gateway 1.12.11.1`	Configures the IPv4 address for the default gateway.

Verifying the Default Gateway Configuration

Use the **show ip route** command to verify the default gateway configuration.

```
switch# show ip route

Codes: C - connected, S - static

Gateway of last resort is 1.12.11.1

S 5.5.5.0/24 via 1.1.1.1, GigabitEthernet1/1
C 1.12.11.0/24 is directly connected, mgmt0
C 1.1.1.0/24 is directly connected, GigabitEthernet1/1
C 3.3.3.0/24 is directly connected, GigabitEthernet1/6
C 3.3.3.0/24 is directly connected, GigabitEthernet1/5
S 3.3.3.0/24 via 1.1.1.1, GigabitEthernet1/1
```

IPv4 Default Network Configuration

If you assign the IPv4 default network address, the switch considers routes to that network as the last resort. If the IPv4 default network address is not available, the switch uses the IPv4 default gateway address. For every network configured with the IPv4 default network address, the switch flags that route as a candidate default route, if the route is available.

Tip If you configure the static route IP forwarding and the default network details, these IPv4 addresses will be used regardless of the default gateway being enabled or disabled. If these IPv4 addresses are configured and not available, the switch will fall back to using the default gateway IPv4 address, if you have configured it. Be sure to configure IPv4 addresses for all entries in the switch if you are using IPv4.

See the "Initial Setup Routine" section on page 5-2 for more information on configuring the IP addresses for all entries in the switch.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IPv4 address of the gateway switch (see Figure 43-2).

Figure 43-2 Overlay VSAN Functionality

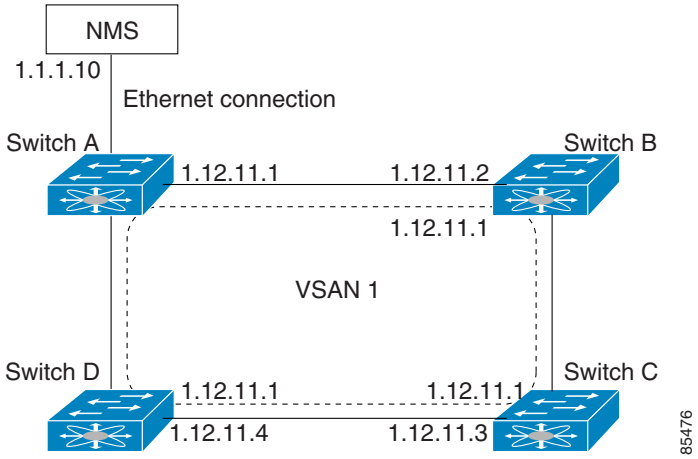

In Figure 43-2, switch A has the IPv4 address 1.12.11.1, switch B has the IPv4 address 1.12.11.2, switch C has the IPv4 address 1.12.11.3, and switch D has the IPv4 address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IPv4 address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IPv4 address (1.12.11.1) in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet world, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface (see the "VSAN Interfaces" section on page 12-40).

To configure default networks using IPv4 addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip default-network 190.10.1.0`	Configures the IPv4 address for the default network (190.10.1.0).
	`switch(config)# ip route 10.0.0.0 255.0.0.0 131.108.3.4` `switch(config)# ip default-network 10.0.0.0`	Defines a static route to network 10.0.0.0 as the static default route.

IPFC

IPFC provides IP forwarding or in-band switch management over a Fibre Channel interface (rather than out-of-band using the Gigabit Ethernet mgmt 0 interface). You can be use IPFC to specify that IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.

Once the VSAN interface is created, you can specify the IP address for that VSAN. You can assign an IPv4 address or an IPv6 address.

Note See the Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces" for information about configuring IPv6 on the Cisco MDS 9000 Family switches.

This topic includes the following sections:

IPFC Configuration Guidelines

Follow these guidelines to configure IPFC:

1. Create the VSAN to use for in-band management, if necessary.
2. Configure an IPv4 address and subnet mask for the VSAN interface.
3. Enable IPv4 routing.
4. Verify connectivity.

Configuring an IPv4 Address in a VSAN

To create a VSAN interface and configure an IPv4 address for that interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures the interface for the specified VSAN (10).
Step 3	`switch(config-if)# ip address 10.0.0.12 255.255.255.0`	Configures the IPv4 address and netmask for the selected interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Verifying the VSAN Interface Configuration

Use the **show interface vsan** command to verify the configuration of the VSAN interface.

Note You can see the output for this command only if you have previously configured a VSAN interface.

```
switch# show interface vsan 1
vsan1 is down (Administratively down)
    WWPN is 10:00:00:0c:85:90:3e:85, FCID not assigned
    Internet address is 10.0.0.12/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
```

Enabling IPv4 Routing

By default, the IPv4 routing feature is disabled in all switches.

To enable the IPv4 routing feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# ip routing`	Enables IPv4 routing (disabled by default).
Step 3	`switch(config)# no ip routing`	Disables IPv4 routing and reverts to the factory settings.

Verifying the IPv4 Routing Configuration

Use the **show ip routing** command to verify the IPv4 routing configuration.

```
switch(config)# show ip routing
ip routing is enabled
```

IPFC Configuration Example

This section describe an example configuration for IPFC. Figure 43-3 shows an example network.

The example network has the following links:

- Switch_1 is connected to the main network by the mgmt 0 interface and to the fabric by an ISL.
- Switch_2 and Switch_3 are connected to the fabric by an ISL but are not connected to the main network.

Figure 43-3 IPFC Example Network

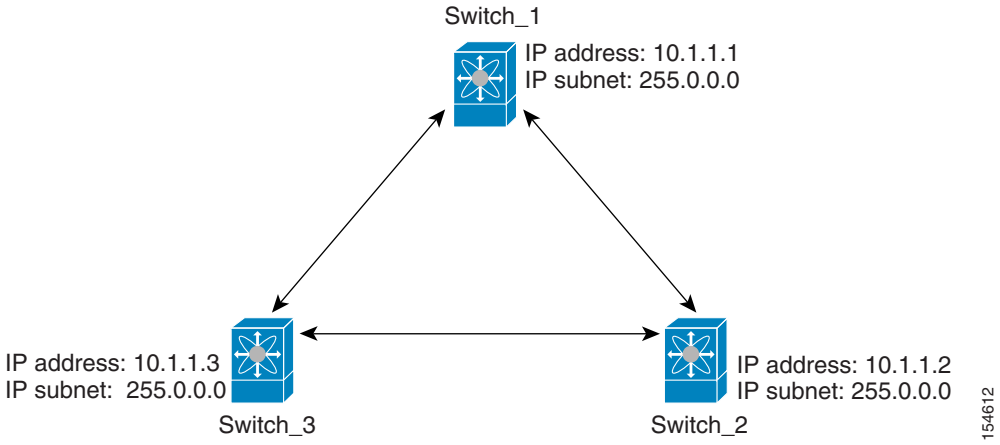

The following steps show how to configure Switch_1 in the example network in Figure 43-3:

Step 1 Create the VSAN interface and enter interface configuration submode.

```
switch_1# config t
switch_1(config)# interface vsan 1
switch_1(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_1(config-if)# ip address 10.1.1.1 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_1(config-if)# no shutdown
switch_1(config-if)# exit
switch_1(config)#
```

Step 4 Enable IPv4 routing.

```
switch_1(config)# ip routing
switch_1(config)# exit
switch_1#
```

Step 5 Display the routes.

```
switch_1# show ip route

Codes: C - connected, S - static

C 172.16.1.0/23 is directly connect, mgmt0
```

```
C 10.0.0.0./8 is directly connected, vsan1
```

The following steps show how to configure Switch_2 in the example network in Figure 43-3.

Step 1 Disable the mgmt 0 interface.

Note Configure this switch using the console connection.

```
switch_2# config t
switch_2(config)# interface mgmt 0
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 2 Create the VSAN interface and enter interface configuration submode.

```
switch_2# config t
switch_2(config)# interface vsan 1
switch_2(config-if)#
```

Step 3 Configure the IP address and subnet mask.

```
switch_2(config-if)# ip address 10.1.1.2 255.0.0.0
```

Step 4 Enable the VSAN interface and exit interface configuration submode.

```
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 5 Enable IPv4 routing.

```
switch_2(config)# ip routing
switch_2(config)# exit
switch_2#
```

Step 6 Display the routes.

```
switch_2# show ip route

Codes: C - connected, S - static

C 10.0.0.0./8 is directly connected, vsan1
```

Step 7 Verify the connectivity to Switch_1.

```
switch_2# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.618 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.528 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.567 ms

--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 4998 ms
rtt min/avg/max/mdev = 0.528/0.570/0.618/0.057 ms
```

The following steps show how to configure Switch_3 in the example network in Figure 43-3.

Step 1 Disable the mgmt 0 interface.

Note Configure this switch using the console connection.

```
switch_3# config t
switch_3(config)# interface mgmt 0
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#

switch_3# config t
switch_3(config)# interface vsan 1
switch_3(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_3(config-if)# ip address 10.1.1.3 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#
```

Step 4 Enable IPv4 routing.

```
switch_3(config)# ip routing
switch_3(config)# exit
switch_3#
```

Step 5 Display the routes.

```
switch_3# show ip route

Codes: C - connected, S - static

C 10.0.0.0./8 is directly connected, vsan1
```

Step 6 Verify the connectivity to Switch_1.

```
switch_3# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=1.19 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.510 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.653 ms

--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2008 ms
rtt min/avg/max/mdev = 0.510/0.787/1.199/0.297 ms
```

IPv4 Static Routes

If your network configuration does not need an external router, you can configure IPv4 static routing on your MDS switch.

Note For information about IPv6 static routing, see the "Configuring IPv6 for Gigabit Ethernet Interfaces" section on page 46-1.

This section includes the following topics:

About IPv4 Static Routes

Static routing is a mechanism to configure IPv4 routes on the switch. You can configure more than one static route.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IPv4 routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

Configuring IPv4 Static Routes

To configure an IPv4 static route, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

Step 2	`switch(config)# ip route <network IP address> <netmask> <next hop IPv4 address> distance <number> interface <vsan number>`
	
	`For example:`
	`switch(config)# ip route 10.0.0.0 255.0.0.0 20.20.20.10 distance 10 interface vsan 1`
	`switch(config)#`
		Configures the static route for the specified IPv4 address, subnet mask, next hop, distance, and interface.

Verifying IPv4 Static Route Information

Use the **show ip route** command to verifying the IPv4 static route configuration.

```
switch# show ip route configured
Destination          Gateway             Mask Metric        Interface

         default      172.22.95.1          0.0.0.0      0            mgmt0
        10.1.1.0          0.0.0.0    255.255.255.0      0            vsan1
     172.22.95.0          0.0.0.0    255.255.255.0      0            mgmt0
```

Use the **show ip route** command to verifying the active and connected IPv4 static route.

```
switch# show ip route

Codes: C - connected, S - static
```

```
Default gateway is 172.22.95.1

C 172.22.95.0/24 is directly connected, mgmt0
C 10.1.1.0/24 is directly connected, vsan1
```

Example 43-1 Displays the IP Routing Status

```
switch# show ip routing
ip routing is disabled
```

Displaying and Clearing ARPs

Address Resolution Protocol (ARP) entries in Cisco MDS 9000 Family switches can be displayed, deleted, or cleared. The ARP feature is enabled on all switches.

- Use the **show arp** command to display the ARP table.

```
switch# show arp
Protocol Address           Age (min)  Hardware Addr  Type  Interface
Internet  171.1.1.1               0  0006.5bec.699c  ARPA  mgmt0
Internet  172.2.0.1               4  0000.0c07.ac01  ARPA  mgmt0
```

- Use the **no arp** command in configuration mode to remove an ARP entry from the ARP table.

```
switch(config)# no arp 172.2.0.1
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
```

Overlay VSANs

This section describes overlay VSANs and how to configure them.

This section includes the following topics:

About Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

Configuring Overlay VSANs

To configure an overlay VSAN, follow these steps:

Step 1 Add the VSAN to the VSAN database on all switch in the fabric.

Step 2 Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side.

Step 3 Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.

Step 4 Configure the default gateway (route) and the IPv4 address on switches that point to the NMS (see Figure 43-4).

Figure 43-4 Overlay VSAN Configuration Example

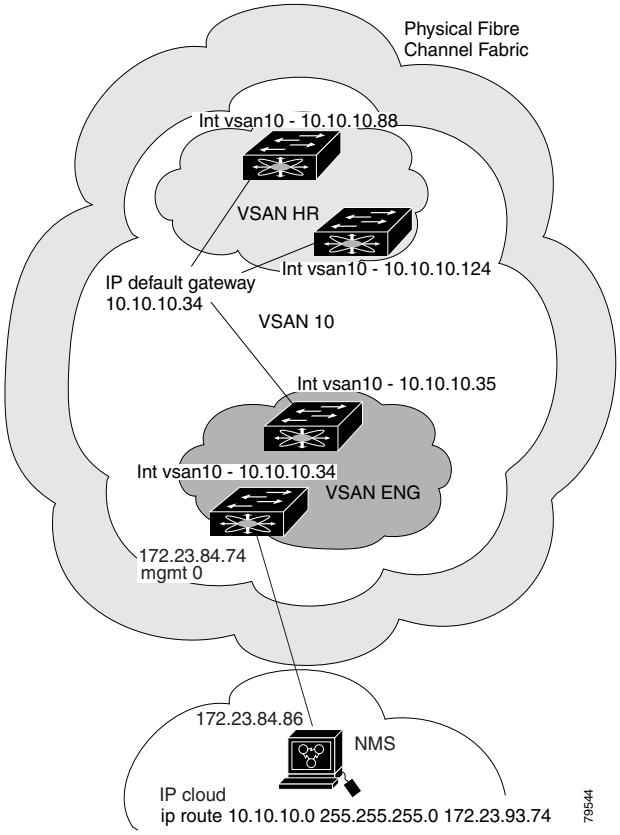

Note To configure the management interface displayed in Figure 43-4, set the default gateway to an IPv4 address on the Ethernet network.

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in Figure 43-4), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch-config-vsan-db#`	Configures the VSAN database.
Step 3	`switch--config-vsan-db# vsan 10 name MGMT_VSAN`	Defines the VSAN in the VSAN database on all of the switches in the Fibre Channel fabric.
Step 4	`switch--config-vsan-db# exit` `switch(config)#`	Exits the VSAN database mode.
Step 5	`switch(config)# interface vsan 10` `switch(config-if)#`	Creates a VSAN interface (VSAN 10).
Step 6	`switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0`	Assigns an IPv4 address and subnet mask for this switch.
Step 7	`switch(config-if)# no shutdown`	Enables the configured interface.
Step 8	`switch(config-if)# end` `switch#`	Exits to EXEC mode.
Step 9	`switch# exit`	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.

To configure the NMS station displayed in Figure 43-4, follow this step:

	Command	Purpose
Step 1	`nms# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.

Multiple VSAN Configuration

More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure multiple VSANs, follow these steps:

Step 1 Add the VSAN to the VSAN database on any switch in the fabric.

Step 2 Create a VSAN interface for the appropriate VSAN on any switch in the fabric.

Step 3 Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.

Step 4 Define the multiple static routes on the Fibre Channel switches and the IP cloud (see Figure 43-5).

Figure 43-5 Multiple VSAN Configuration Example

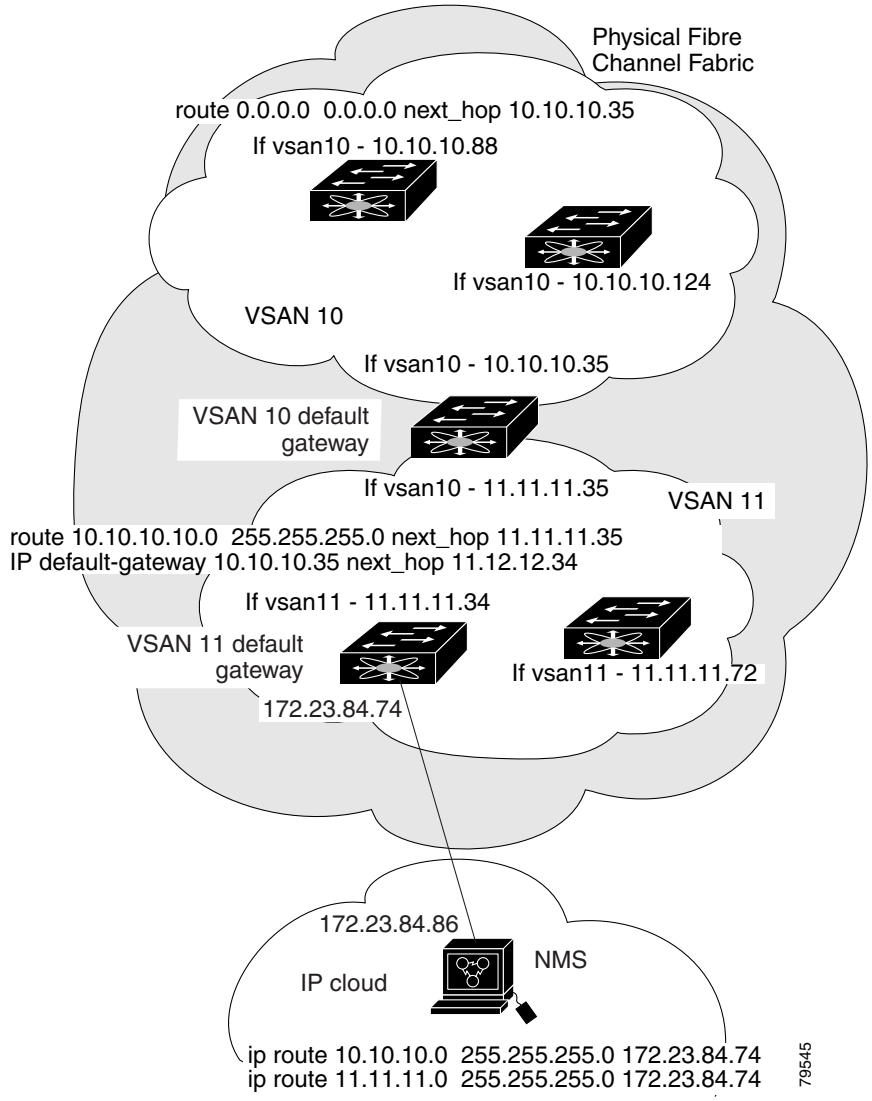

To configure an overlay VSAN (using the example in Figure 43-5), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# vsan database` `switch-config-vsan-db#`	Configures the VSAN database.
Step 3	`switch-config-vsan-db# vsan 10 name MGMT_VSAN_10` `switch-config-vsan-db#`	Defines the VSAN in the VSAN database on all of the switches in VSAN 10.
Step 4	`switch-config-vsan-db# exit` `switch(config)#`	Exits the VSAN database configuration submode.
Step 5	`switch-config-vsan-db# vsan 11 name MGMT_VSAN_11` `switch-config-vsan-db#`	Defines the VSAN in the VSAN database on all of the switches in VSAN 11.

	Command	Purpose
Step 6	`switch-config-vsan-db# exit` `switch(config)#`	Exits the VSAN database configuration submode.
Step 7	`switch(config)# interface vsan 10` `switch(config-if)#`	Enters the interface configuration submode for VSAN 10.
Step 8	`switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0` `switch(config-if)#`	Assigns an IPv4 address and subnet mask for this interface.
Step 9	`switch(config-if)# no shutdown`	Enables the configured interface for VSAN 10.
Step 10	`switch(config-if)# exit` `switch(config)#`	Exits the VSAN 10 interface mode.
Step 11	`switch(config)# interface vsan 11` `switch(config-if)#`	Enters the interface configuration submode for VSAN 11.
Step 12	`switch(config-if)# ip address 11.11.11.0 netmask 255.255.255.0` `switch(config-if)#`	Assigns an IPv4 address and subnet mask for this interface.
Step 13	`switch(config-if)# no shutdown`	Enables the configured interface for VSAN 11.
Step 14	`switch(config-if)# end` `switch#`	Exits to EXEC mode.
Step 15	`switch# exit`	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.
Step 16	`NMS# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the IPv4 cloud.
Step 17	`NMS# route ADD 11.11.11.0 MASK 255.255.255.0 172.22.93.74`	Defines a static route for VSAN 11 on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.
Step 18	`switch# route 10.10.10.0 255.255.255.0 next_hop 11.11.11.35`	Defines the route to reach subnet 10 from subnet 11.

Virtual Router Redundancy Protocol

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. This section provides details on the VRRP feature.

This section includes the following topics:

About VRRP

VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and the draft-ietf-vrrp-ipv6 specification.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- Interface Mgmt 0 supports only one VRRP group. All other interface supports up to 7 virtual router groups, including both IPv4 and IPv6 combined.
- VR IDs can be reused in multiple VSANs with different virtual router IP mapping.
- Both IPv4 and IPv6 is supported.
- The management interface (mgmt 0) supports only one virtual router group. All other interfaces each support up to seven virtual router groups, including both IPv4 and IPv6 combined. Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.

Note If you are using IPv6, you must either configure an IPv6 address on the interface or enable IPv6 on the interface. For more information about IPv6, see Chapter 46, "Configuring IPv6 for Gigabit Ethernet Interfaces."

In Figure 43-6, switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have an IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches do not have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 43-6 VRRP Functionality

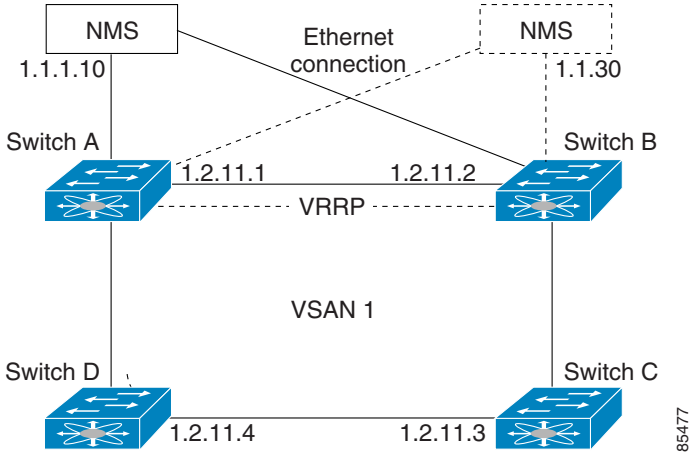

In Figure 43-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 43-7 Redundant Gateway

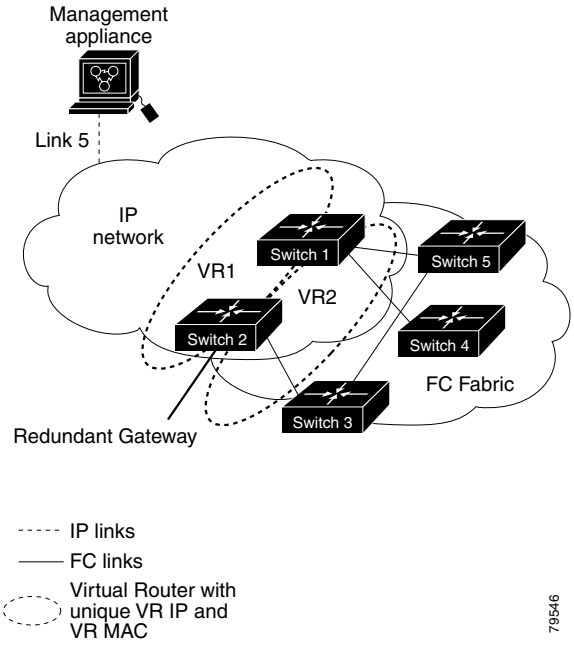

Configuring VRRP

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

Adding and Deleting Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

Note The total number of VRRP groups that you can configure on a Gigabit Ethernet port, including main interfaces and subinterfaces, cannot exceed seven. This limitation applies to both IPv4 and IPv6 groups.

To create or remove a VR for IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates VR ID 250.
	`switch(config-if)# no vrrp 250`	Removes VR ID 250.

To create or remove a VR for IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp ipv6 250` `switch(config-if-vrrp-ipv6)#`	Creates VR ID 250.
	`switch(config-if)# no vrrp ipv6 250`	Removes VR ID 250.

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address, either IPv4 or IPv6, before attempting to enable a VR.

To enable or disable a virtual router configure for IPv4, follow these steps:

	Command	Purpose
Step 1	`switch(config-if-vrrp)# no shutdown`	Enables VRRP configuration.
	`switch(config-if-vrrp)# shutdown`	Disables VRRP configuration.

To enable or disable a virtual router configured for IPv6, follow these steps:

	Command	Purpose
Step 1	`switch(config-if-vrrp-ipv6)# no shutdown`	Enables VRRP configuration.
	`switch(config-if-vrrp-ipv6)# shutdown`	Disables VRRP configuration.

Adding Virtual Router IP Addresses

One virtual router IP address can be configured for a virtual router. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address. You can configure either an IPv4 address or an IPv6 address.

According to the VRRP specification, the master VRRP router drops the packets addressed to the virtual router's IP address because the virtual router is only intended as a next-hop router to forward packets. In MDS switches however, some applications require that packets addressed to virtual router's IP address be accepted and delivered to them. By using the **secondary** option to the virtual router IPv4 address, the VRRP router will accept these packets when it is the master.

To configure an IPv4 address for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# interface ip address 10.0.0.12 255.255.255.0`	Configures an IPv4 address and subnet mask. The IPv4 address must be configured before the VRRP is added.
Step 4	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates VR ID 250.
Step 5	`switch(config-if-vrrp)# address 10.0.0.10`	Configures the IPv4 address for the selected VR. **Note** This IP v4address should be in the same subnet as the IPv4 address of the interface.
	`switch(config-if-vrrp)# no address 10.0.0.10`	Removes the IP address for the selected VR.
Step 6	`switch(config-if-vrrp)# address 10.0.0.10 secondary`	Configures the IP address (10.0.0.10) as secondary for the selected VR. **Note** The **secondary** option should be used only with applications that require VRRP routers to accept the packets sent to the virtual router's IP address and deliver to them. For example, iSNS requires this option (see the "Enabling the iSNS Server" section on page 42-88).
	`switch(config-if-vrrp)# no address 10.0.0.10 secondary`	Removes the IP address (10.0.0.10) as secondary for the selected VR.

To configure an IPv6 address for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 12` `switch(config-if)#`	Configures a VSAN interface (VSAN 12).
Step 3	`switch(config-if)# interface ipv6 address 2001:0db8:800:200c::417a/64`	Configures an IP address and prefix. The IPv6 address must be configured before the VRRP is added.
Step 4	`switch(config-if)# vrrp ipv6 200` `switch(config-if-vrrp-ipv6)#`	Creates VR ID 200.

	Command	Purpose
Step 5	`switch(config-if-vrrp-ipv6)# address 2001:0db8:800:200c::417a`	Assigns single primary link-local IPv6 address or one of the multiple secondary IPv6 addresses. **Note** If this IPv6 address is the same as the physical IPv6 address, this switch is automatically the owner of this IPv6 address.
	`switch(config-if-vrrp-ipv6)# no address 2001:0db8:800:200c::417a`	Removes the IPv6 address for the selected VR.

Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for switches with the primary IP address.

To set the priority for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# priority 2`	Configures the priority for the selected VRRP. **Note** Priority 255 cannot be preempted.
	`switch(config-if-vrrp)# no priority`	Reverts to the default value (100 for switch with the secondary IPv4 addresses and 255 for switches with the primary IPv4 address).

To set the priority for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 12` `switch(config-if)#`	Configures a VSAN interface (VSAN 12).
Step 3	`switch(config-if)# vrrp ipv6 200` `switch(config-if-vrrp-ipv6)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp-ipv6)# priority 2`	Configures the priority for the selected VRRP. **Note** Priority 255 cannot be preempted.
	`switch(config-if-vrrp-ipv6)# no priority`	Reverts to the default value (100 for switch with the secondary IPv6 addresses and 255 for switches with the primary IPv6 address).

Time Interval for Advertisement Packets

The valid time range for an advertisement packet on an interface using IPv4 is between 1 and 255 seconds. The default value is 1 (one) second. If the switch has the primary IP address, this time must be specified.

To set the time interval for advertisement packets for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp 50` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# advertisement-interval 15`	Sets the interval time in seconds between sending advertisement frames. The range is 1 to 255.
	`switch(config-if-vrrp)# no advertisement-interval`	Reverts to the default value (1 second).

To set the time interval for advertisement packets for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 12` `switch(config-if)#`	Configures a VSAN interface (VSAN 12).
Step 3	`switch(config-if)# vrrp ipv6 200` `switch(config-if-vrrp-ipv6)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp-ipv6)#` `advertisement-interval 150`	Sets the interval time in centiseconds between sending advertisement frames. The range is 100 to 4095. The default is 100 centiseconds.
	`switch(config-if-vrrp-ipv6)# no` `advertisement-interval`	Reverts to the default value (100 centiseconds).

Priority Preemption

You can enable a higher priority backup virtual router to preempt the lower priority master virtual router.

Note If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Note The VRRP preemption is not supported on IP storage Gigabit Ethernet interfaces.

To enable or disable preempting when using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# preempt`	Enables the higher priority backup virtual router to preempt the lower priority master virtual router. **Note** This preemption does not apply to the primary IP address.
	`switch(config-if-vrrp)# no preempt`	Disables (default) the preempt option and allows the master to keep its priority level.

To enable or disable preempting when using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 12` `switch(config-if)#`	Configures a VSAN interface (VSAN 12).
Step 3	`switch(config-if)# vrrp ipv6 200` `switch(config-if-vrrp-ipv6)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp-ipv6)# preempt`	Enables the higher priority backup virtual router to preempt the lower priority master virtual router. **Note** This preemption does not apply to the primary IP address.
	`switch(config-if-vrrp-ipv6)# no preempt`	Disables (default) the preempt option and allows the master to keep its priority level.

Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.

Note All VRRP configurations must be duplicated.

Note VRRP router authentication does not apply to IPv6.

To set an authentication option for a virtual router, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 1` `switch(config-if)#`	Configures a VSAN interface (VSAN 1).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# authentication text password`	Assigns the simple text authentication option and specifies the password for this option.
	`switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003`	Assigns the MD5 authentication option and specifies the key and the unique SPI value for this option. The SPI and the valid range is 0x100 to 0xFFFFFFFF.
	`switch(config-if-vrrp)# no authentication`	Assigns the no authentication option, which is the default.

Priority Based on Interface State Tracking

Interface state tracking changes the priority of the virtual router based on the state of another interface in the switch. When the tracked interface is down, the priority reverts to the priority value for the virtual router (see the"Priority for the Virtual Router" section on page 43-21). When the tracked interface is up, the priority of the virtual router is restored to the interface state tracking value. You can track the state of either a specified VSAN interface or the management interface (mgmt 0). The interface state tracking feature is disabled by default.

Note For interface state tracking to function, you must enable preemption on the interface. See the "Priority Preemption" section on page 43-22.

To track the interface priority for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 10` `switch(config-if)#`	Configures a VSAN interface (VSAN 10).
Step 3	`switch(config-if)# vrrp 250` `switch(config-if-vrrp)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp)# preempt`	Enables priority preemption.
Step 5	`switch(config-if-vrrp)# track interface mgmt 0 priority 2`	Specifies the priority of the virtual router to be modified based on the state of the management interface.
	`switch(config-if-vrrp)# no track`	Disables the tracking feature.

To track the interface priority for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 12` `switch(config-if)#`	Configures a VSAN interface (VSAN 12).
Step 3	`switch(config-if)# vrrp ipv6 200` `switch(config-if-vrrp-ipv6)#`	Creates a virtual router.
Step 4	`switch(config-if-vrrp-ipv6)# preempt`	Enables priority preemption.
Step 5	`switch(config-if-vrrp-ipv6)# track interface mgmt 0 priority 2`	Specifies the priority of the virtual router to be modified based on the state of the management interface. **Note** You must enable IPv6 on the tracked interface for the priority tracking to take affect (see the "Configuring Basic Connectivity for IPv6" section on page 46-11). If IPv6 is not enabled, the interface state is treated as down by VRRP over IPv6, regardless of the actual state of the interface.
	`switch(config-if-vrrp-ipv6)# no track`	Disables the tracking feature.

Displaying IPv4 VRRP Information

Use the **show vrrp vr** command to display configured IPv4 VRRP information (see Examples 43-2 to 43-4).

Example 43-2 Displays IPv4 VRRP Configured Information

```
switch# show vrrp vr 7 interface vsan 2 configuration
vr id 7 configuration
admin state down
priority 100
no authentication
advertisement-Interval 1
preempt yes
tracking interface vsan1 priority 2
protocol IP
```

Example 43-3 Displays IPv4 VRRP Status Information

```
switch# show vrrp vr 7 interface vsan 2 status
vr id 7 status
MAC address 00:00:5e:00:01:07
Operational state: init
```

Example 43-4 Displays IPv4 VRRP Statistics

```
switch# show vrrp vr 7 interface vsan 2 statistics
vr id 7 statistics
Become master 0
Advertisement 0
Advertisement Interval Error 0
```

```
Authentication Failure 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Authentication Type 0
Mismatch Authentication 0
Invalid Packet Length 0
```

Displaying IPv6 VRRP Information

Use the **show vrrp ipv6 vr** command to display configured IPv6 VRRP information (see Example 43-5 through Example 43-9).

Example 43-5 Displays IPv6 VRRP Information

```
switch# show vrrp ipv6 vr 1
      Interface  VR IpVersion Pri   Time Pre State   VR IP addr
---------------------------------------------------------------------------
        GigE1/5   1   IPv6    100  100cs     master  2004::1
        GigE1/6   1   IPv6    100  100cs     backup  2004::1
```

Example 43-6 Displays IPv6 VRRP Interface Configuration Information

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 configuration
IPv6 vr id 1 configuration
admin state up
priority 100
associated ip: 2004::1
advertisement-interval 100
preempt no
protocol IPv6
```

Example 43-7 Displays IPv6 VRRP Interface Status Information

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 status
IPv6 vr id 1 status
MAC address 00:00:5e:00:02:01
Operational state: master
Up time 37 min,  10 sec
Master IP address: fe80::20c:30ff:feda:96dc
```

Example 43-8 Displays IPv6 VRRP Statistics

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 statistics
IPv6 vr id 1 statistics
Become master 1
Advertisement 0
Advertisement Interval Error 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Packet Lenght 0
```

Displaying VRRP Statistics

Use the **show vrrp statistics** command to display configured IPv6 VRRP information (see Example 43-9).

Example 43-9 Displays VRRP Cumulative Statistics

```
switch# show vrrp statistics
Invalid checksum 0
Invalid version 0
Invalid VR ID 0
```

Clearing VRRP Statistics

Use the **clear vrrp statistics** command to clear all the VRRP statistics for all interfaces on the switch (see Example 43-10).

Example 43-10 Clears VRRP Statistics

```
switch# clear vrrp Statistics
```

Use the **clear vrrp vr** command to clear both the IPv4 and IPv6 VRRP statistics for a specified interface (see Example 43-10).

Example 43-11 Clears VRRP Statistics on a Specified Interface

```
switch# clear vrrp vr 1 interface vsan 1
```

Use the **clear vrrp ipv4** command to clear all the statistics for the specified IPv4 virtual router (see Example 43-12).

Example 43-12 Clears VRRP IPv4 Statistics on a Specified Interface

```
switch# clear vrrp ipv4 vr 7 interface vsan 2
```

Use the **clear vrrp ipv6** command to clear all the statistics for the specified IPv6virtual router (see Example 43-13).

Example 43-13 Clears VRRP IPv6 Statistics on a Specified Interface

```
switch# clear vrrp ipv6 vr 7 interface vsan 2
```

DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts because of one of the following reasons:

- The IP address or the switch name is wrongly configured.
- The DNS server is not reachable because external reasons (reasons beyond our control).

Note When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

To configure a DNS server, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ip domain-lookup`	Enables the IP Domain Naming System (DNS)-based host name-to-address translation.
	`switch(config)# no ip domain-lookup`	Disables (default) the IP DNS-based host name-to-address translation and reverts to the factory default.
Step 3	`switch(config)# ip domain-name cisco.com`	Enables the default domain name feature used to complete unqualified host names. Any IP host name that does not contain a domain name (that is, any name without a dot) will have the dot and cisco.com appended to it before being added to the host table.
	`switch(config)# no ip domain-name cisco.com`	Disables (default) the domain name.

Step 4	`switch(config)# ip domain-list harvard.edu`
	`switch(config)# ip domain-list stanford.edu`
	`switch(config)# ip domain-list yale.edu`
		Defines a filter of default domain names to complete unqualified host names by using the **ip domain-list** global configuration command. You can define up to 10 domain names in this filter. To delete a name from a filter, use the **no** form of this command.

	`switch(config)# no ip domain-list`	Deletes the defined filter and reverts to factory default. No domains are configured by default.
	Note If you have not configured a domain list, the domain name that you specified with the **ip domain-name** global configuration command is used. If you configured a domain list, the default domain name is not used. The **ip domain-list** command is similar to the **ip domain-name** command, except that with the **ip domain-list** command you can define a list of domains, each to be tried in turn.	
Step 5	`switch(config)# ip name-server 15.1.0.1 2001:0db8:800:200c::417a`	Specifies the first address (15.1.0.1) as the primary server and the second address (2001:0db8:800:200c::417a) as the secondary server. You can configure a maximum of six servers.
	`switch(config)# no ip name-server`	Deletes the configured server(s) and reverts to factory default. No server is configured by default.
Step 6	**Note** Alternatively, you can configure the DNS entry using the switch names (instead of IP addresses). The configured switch name automatically looks up the corresponding IP address.	

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see Example 43-14).

Example 43-14 Displays Configured Host Details

```
switch# show hosts
Default domain is cisco.com
Domain list: ucsc.edu harvard.edu yale.edu stanford.edu
Name/address lookup uses domain service
Name servers are 15.1.0.1 15.2.0.0
```

Default Settings

Table 43-1lists the default settings for DNS features.

Table 43-1 Default DNS Settings

Parameters	Default
Domain lookup	Disabled.
Domain name	Disabled.
Domains	None.
Domain server	None.
Maximum domain servers	6.

Table 43-2lists the default settings for VRRP features.

Table 43-2 Default VRRP Settings

Parameters	Default
Virtual router state	Disabled.
Maximum groups per VSAN	255.
Maximum groups per Gigabit Ethernet port	7.
Priority preemption	Disabled.
Virtual router priority	100 for switch with secondary IP addresses. 255 for switches with the primary IP address.
Priority interface state tracking	Disabled.
Advertisement interval	1 second for IPv4. 100 centiseconds for IPv6.

CHAPTER 44

Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

Note FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

This chapter includes the following sections:

Services Modules

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types of storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2).

Gigabit Ethernet ports in these modules can be configured to support the FCIP protocol, the iSCSI protocol, or both protocols simultaneously.:

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. Figure 44-1 shows how the IPS module is used in different FCIP scenarios.

Figure 44-1 FCIP Scenarios

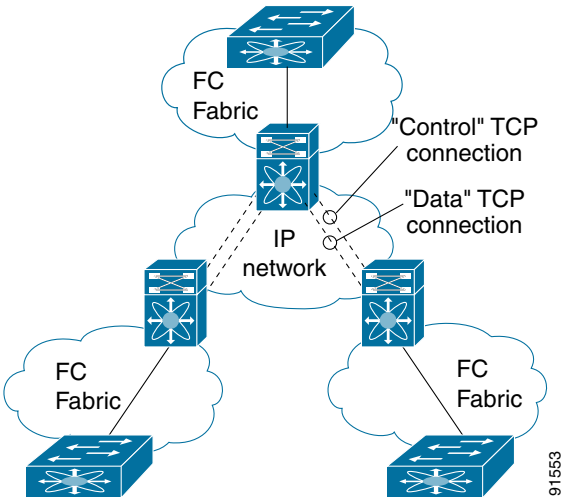

- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. Figure 44-2 depicts the iSCSI scenarios in which the IPS module is used.

Figure 44-2 iSCSI Scenarios

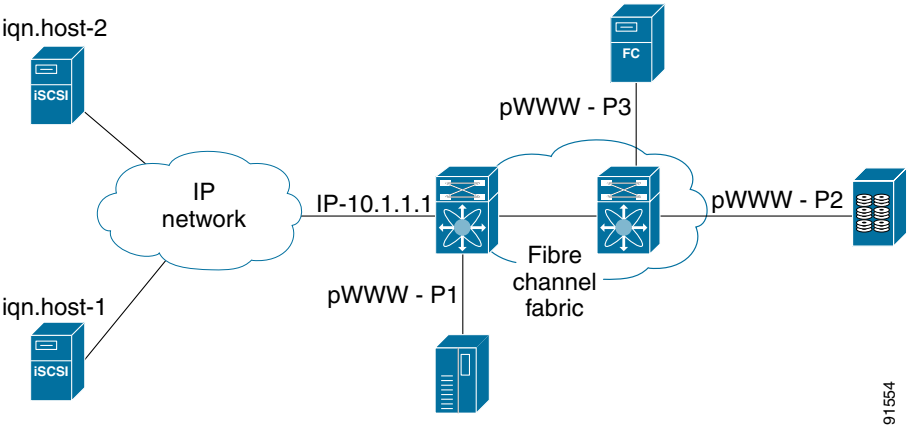

Module Status Verification

After inserting the module, verify the status of the module using the **show module** command:

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
1    0      Caching Services Module           DS-X9560-SMAP      ok
2    8      IP Storage Services Module        DS-X9308-SMIP      ok <-------IPS-8 module
4    16     2x1GE IPS, 14x1/2Gbps FC Module   DS-X9216i-K9-SUP   ok <-------MPS-14/2 module
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
9    4      IP Storage Services Module        DS-X9304-SMIP      ok <---------IPS-4 module

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
1    2.0(1)       0.201   20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
2    2.0(1)       0.201   20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
4    2.0(1)       0.201   20:c1:00:05:30:00:07:1e to 20:d0:00:05:30:00:07:1e
5    2.0(1)       0.0     --
6    2.0(1)       0.0     --
9    2.0(1)       0.1     22:01:00:05:30:00:07:1e to 22:04:00:05:30:00:07:1e

Mod      Application Image Description       Application Image Version
-------- ------------------------------      -------------------------
1        svc-node1                           1.3(5M)
1        svc-node2                           1.3(5M)

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    00-05-30-01-49-c2 to 00-05-30-01-4a-46  JAB073907EP
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de  JAB064605a2
4    00-05-30-01-7f-32 to 00-05-30-01-7f-38  JAB081405AM
5    00-05-30-00-2c-4e to 00-05-30-00-2c-52  JAB06350B1M
6    00-05-30-00-19-66 to 00-05-30-00-19-6a  JAB073705GL
9    00-0d-bc-2f-d6-00 to 00-0d-bc-2f-d6-08  JAB080804TN

* this terminal session
```

IPS Module Upgrade

Caution A software upgrade is only disruptive for the IPS module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

MPS-14/2 Module Upgrade

Caution A software upgrade is only partially disruptive for the MPS-14/2 module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and 2 Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one or more of the following hardware:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).

Note In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel ports and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

IPS Module Core Dumps

IPS core dumps are different from the system's kernel core dumps for other modules. When the IPS module's operating system (OS) unexpectedly resets, it is useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

  Use the **show cores** command to list these files.

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module because of its large space requirement. They are copied directly to an external TFTP server.

  Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

To configure IPS core dumps on the IPS module, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ips core dump full` `ips core dump full' successfully set for module 9`	Configures a dump of the full core generation for all IPS modules in the switch.
	`switch(config)# no ips core dump full` `ips core dump partial' successfully set for module 9`	Configures a dump of the partial core (default) generation for the IPS module in slot 9.

Configuring Gigabit Ethernet High Availability

Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 44-3).

Figure 44-3 VRRP Scenario

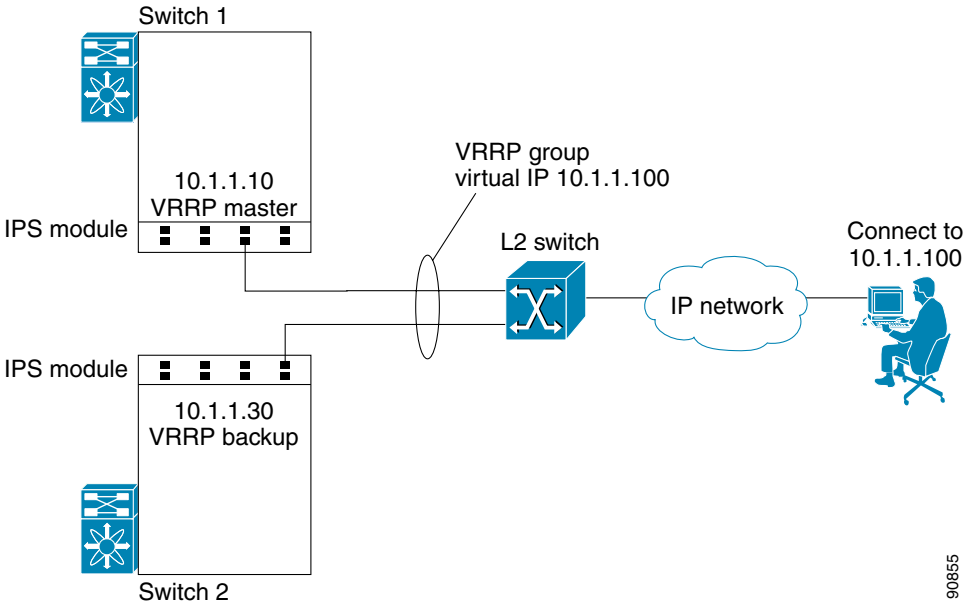

In Figure 44-3, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch

- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces

See the "Virtual Router Redundancy Protocol" section on page 43-16.

Note You can configure no more than seven VRRP groups, both IPv4 and IPv6, on a Gigabit Ethernet interface, including the main interface and all subinterfaces.

Configuring VRRP for Gigabit Ethernet Interfaces

To configure VRRP for Gigabit Ethernet interfaces using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch1(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# ip address 10.1.1.10 255.255.255.0`	Assigns the IPv4 address (10.1.1.10) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	`switch(config-if)# no shutdown`	Enables the selected interface.
Step 5	`switch(config-if)# vrrp 100` `switch(config-if-vrrp)`	Creates VR ID 100.
Step 6	`switch(config-if-vrrp)# address 10.1.1.100`	Configures the virtual IPv4 address (10.1.1.100) for the selected VRRP group (identified by the VR ID). **Note** The virtual IPv4 address must be in the same subnet as the IPv4 address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IPv4 address.
Step 7	`switch(config-if-vrrp)# priority 10`	Configures the priority for the selected interface within this VRRP group. **Note** The interface with the highest priority is selected as the master.
Step 8	`switch(config-if-vrrp)# no shutdown`	Enables the VRRP protocol on the selected interface.

To configure VRRP for Gigabit Ethernet interfaces using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch1(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# ipv6 address 2001:0db8:800:200c::417a/64`	Assigns the IPv6 address for the Gigabit Ethernet interface.

	Command	Purpose
Step 4	`switch(config-if)# no shutdown`	Enables the selected interface.
Step 5	`switch(config-if)# vrrp ipv6 100` `switch(config-if-vrrp-ipv6)`	Creates VR ID 100.
Step 6	`switch(config-if-vrrp-ipv6)# address 2001:0db8:800:200c::417a`	Assigns single primary link-local IPv6 address or one of the multiple secondary IPv6 addresses. **Note** If this IPv6 address is the same as the physical IPv6 address, this switch is automatically the owner of this IPv6 address.
Step 7	`switch(config-if-vrrp-ipv6)# priority 10`	Configures the priority for the selected interface within this VRRP group. **Note** The interface with the highest priority is selected as the master.
Step 8	`switch(config-if-vrrp-ipv6)# no shutdown`	Enables the VRRP protocol on the selected interface.

Note If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

Note The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

About Ethernet PortChannel Aggregation

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that FCIP link.

Note The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3ad protocol.

Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see Figure 44-4).

Note PortChannel members must be one of these combinations: ports 1–2, ports 3–4, ports 5–6, or ports 7–8.

Figure 44-4 Ethernet PortChannel Scenario

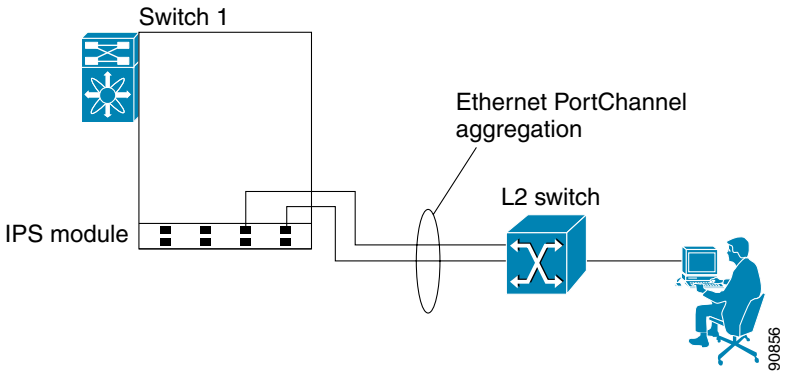

In Figure 44-4, Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.

Note PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

Configuring Ethernet PortChannels

The PortChannel configuration specified in Chapter 16, "Configuring PortChannels," also applies to Ethernet PortChannel configurations.

To configure Ethernet PortChannels, follow these steps:

	Command	Purpose
Step 1	`switch1# config terminal` `switch1(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface port-channel 10` `switch(config-if)#`	Configures the specified PortChannel (10).
Step 3	`switch(config-if)# ip address 10.1.1.1 255.255.255.0`	Enters the IPv4 address (10.1.1.1) and subnet mask (255.255.255.0) for the PortChannel. **Note** A PortChannel does not have any members when first created.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.
Step 5	`switch(config)# interface gigabitethernet 9/3` `switch(config-if)#`	Configures the specified Gigabit Ethernet interface (slot 9, port 3).

	Command	Purpose

Step 6	`switch(config-if)# channel-group 10`
	`gigabitethernet 9/3 added to port-channel 10 and disabled`
	`please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up`
	`switch(config-if)#`
		Adds Gigabit Ethernet interfaces 9/3 to channel group 10. If channel group 10 does not exist, it is created. The port is shut down.

Step 7	`switch(config-if)# no shutdown`	Enables the selected interface.
Step 8	`switch(config)# interface gigabitethernet 9/4` `switch(config-if)#`	Configures the specified Gigabit Ethernet interface (slot 9, port 4).

Step 9	`switch(config-if)# channel-group 10`
	`gigabitethernet 9/4 added to port-channel 10 and disabled`
	`please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up`
		Adds Gigabit Ethernet interfaces 9/4 to channel group 10. The port is shut down.

Step 10	`switch(config-if)# no shutdown`	Enables the selected interface.

Note Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:
- The interface already has an IP address assigned.
- The subinterfaces are configured on that interface.
- The interface already has an associated IPv4-ACL rule and the PortChannel does not.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS module or MPS-14/2 module.

See the "Configuring CDP" section on page 5-36.

Displaying Statistics

This section provides examples to verify Gigabit Ethernet and TCP/IP statistics on the IP storage ports.

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface gigabitethernet** command on each switch to verify that the interfaces are up and functioning as desired. See Example 44-1 and Example 44-2.

Example 44-1 Displays the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
GigabitEthernet8/1 is up              <----------The interface is in the up state.
    Hardware is GigabitEthernet, address is 0005.3000.a98e
    Internet address is 10.1.3.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    Port mode is IPS
```

```
    Speed is 1 Gbps
    Beacon is turned off
    5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    3343 packets input, 406582 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    8 packets output, 336 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Example 44-2 Displays the Gigabit Ethernet Subinterface

```
switch# show interface gigabitethernet 4/2.100
GigabitEthernet4/2.100 is up
    Hardware is GigabitEthernet, address is 0005.3000.abcb
    Internet address is 10.1.2.100/24
    MTU 1500 bytes
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 packets input, 0 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    1 packets output, 46 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Displaying Ethernet MAC Statistics

The **show ips stats mac interface gigabitethernet** command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See Example 44-3.

Note Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

Example 44-3 Displays Ethernet MAC Statistics

```
switch# show ips stats mac interface gigabitethernet 8/1
Ethernet MAC statistics for port GigabitEthernet8/1
    Hardware Transmit Counters
      237 frame 43564 bytes
      0 collisions, 0 late collisions, 0 excess collisions
      0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
    Hardware Receive Counters
      427916 bytes, 3464 frames, 0 multicasts, 3275 broadcasts
      0 bad, 0 runt, 0 CRC error, 0 length error
      0 code error, 0 align error, 0 oversize error
    Software Counters
      3429 received frames, 237 transmit frames
      0 frames soft queued, 0 current queue, 0 max queue
      0 dropped, 0 low memory
```

Displaying DMA-Bridge Statistics

You can display direct memory access (DMA) device statistics using the **show ips stats dma-bridge interface gigabitethernet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns DMA bridge statistics for that interface. See Example 44-4.

Note Use the physical interface, not the subinterface, to display DMA-bridge statistics.

Example 44-4 Displays DMA-Bridge Statistics

```
switch# show ips stats dma-bridge interface gigabitethernet 7/1
Dma-bridge ASIC Statistics for port GigabitEthernet7/1
    Hardware Egress Counters
      231117 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
    Hardware Ingress Counters
      218255 Good, 0 protocol error, 0 header checksum error
      0 FC CRC error, 0 iSCSI CRC error, 0 parity error
    Software Egress Counters
      231117 good frames, 0 bad header cksum, 0 bad FIFO SOP
      0 parity error, 0 FC CRC error, 0 timestamp expired error
      0 unregistered port index, 0 unknown internal type
      0 RDL ok, 0 RDL drop (too big), 0 RDL ttl_1
      3656368645 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
      Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
    Software Ingress Counters
      218255 Good frames, 0 header cksum error, 0 FC CRC error
      0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
      0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
      0 out of memory drop, 0 queue full drop
      0 RDL ok, 0 RDL drop (too big)
      Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
```

This output shows all Fibre Channel frames that ingress or egress from the Gigabit Ethernet port.

Displaying TCP Statistics

Use the **show ips stats tcp interface gigabitethernet** to display and verify TCP statistics. This command takes the main Ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The **detail** option shows all information maintained by the interface. See Example 44-5 and Example 44-6.

Example 44-5 Displays TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
    Connection Stats
      0 active openings, 3 accepts
      0 failed attempts, 12 reset received, 3 established
    Segment stats
      163 received, 355 sent, 0 retransmitted
      0 bad segments received, 0 reset sent
    TCP Active Connections
      Local  Address        Remote Address        State      Send-Q   Recv-Q
      0.0.0.0:3260          0.0.0.0:0             LISTEN     0        0
```

Example 44-6 Displays Detailed TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1 detail
TCP Statistics for port GigabitEthernet4/1
    TCP send stats
      355 segments, 37760 bytes
      222 data, 130 ack only packets
      3 control (SYN/FIN/RST), 0 probes, 0 window updates
      0 segments retransmitted, 0 bytes
      0 retransmitted while on ethernet send queue, 0 packets split
      0 delayed acks sent
    TCP receive stats
      163 segments, 114 data packets in sequence, 6512 bytes in sequence
      0 predicted ack, 10 predicted data
      0 bad checksum, 0 multi/broadcast, 0 bad offset
      0 no memory drops, 0 short segments
      0 duplicate bytes, 0 duplicate packets
      0 partial duplicate bytes, 0 partial duplicate packets
      0 out-of-order bytes, 1 out-of-order packets
      0 packet after window, 0 bytes after window
      0 packets after close
      121 acks, 37764 ack bytes, 0 ack toomuch, 4 duplicate acks
      0 ack packets left of snd_una, 0 non-4 byte aligned packets
      8 window updates, 0 window probe
      30 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
    TCP Connection Stats
      0 attempts, 3 accepts, 3 established
      3 closed, 2 drops, 0 conn drops
      0 drop in retransmit timeout, 1 drop in keepalive timeout
      0 drop in persist drops, 0 connections drained
    TCP Miscellaneous Stats
      115 segments timed, 121 rtt updated
      0 retransmit timeout, 0 persist timeout
      12 keepalive timeout, 11 keepalive probes
    TCP SACK Stats
      0 recovery episodes, 0 data packets, 0 data bytes
      0 data packets retransmitted, 0 data bytes retransmitted
      0 connections closed, 0 retransmit timeouts
    TCP SYN Cache Stats
      15 entries, 3 connections completed, 0 entries timed out
      0 dropped due to overflow, 12 dropped due to RST
      0 dropped due to ICMP unreach, 0 dropped due to bucket overflow
      0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
      0 hash collisions, 0 retransmitted
    TCP Active Connections
      Local  Address        Remote Address        State      Send-Q   Recv-Q
      0.0.0.0:3260          0.0.0.0:0             LISTEN     0        0
```

Use the **show ips stats icmp interface gigabitethernet** to display and verify IP statistics. This command takes the main Ethernet interface as a parameter and returns the ICMP statistics for that interface. See Example 44-7.

Example 44-7 Displays ICMP Statistics

```
switch# show ips stats icmp interface gigabitethernet 2/1
ICMP Statistics for port GigabitEthernet2/1
    0 ICMP messages received
    0 ICMP messages dropped due to errors
    ICMP input histogram
      0 destination unreachable
      0 time exceeded
      0 parameter problem
```

```
    0 source quench
    0 redirect
    0 echo request
    0 echo reply
    0 timestamp request
    0 timestamp reply
    0 address mask request
    0 address mask reply
  ICMP output histogram
    0 destination unreachable
    0 time exceeded
    0 parameter problem
    0 source quench
    0 redirect
    0 echo request
    0 echo reply
    0 timestamp request
    0 timestamp reply
    0 address mask request
    0 address mask reply
```

Default Settings

Table 44-1 lists the default settings for IP storage services parameters.

Table 44-1 Default Gigabit Ethernet Parameters

Parameters	Default
IPS core size	Partial

CHAPTER 45

Configuring IPv4 for Gigabit Ethernet Interfaces

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. This chapter describes how to configure IPv4 addresses and other IPv4 features.

This chapter includes the following topics:

About IPv4

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

Note For information about configuring FCIP, see Chapter 40, "Configuring FCIP." For information about configuring iSCSI, see Chapter 42, "Configuring iSCSI."

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.

Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.

Note To configure IPv6 on a Gigabit Ethernet interface, see the "Configuring IPv6 Addressing and Enabling IPv6 Routing" section on page 46-11.

Tip Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration for IPv4

Figure 45-1 shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 45-1 Gigabit Ethernet IPv4 Configuration Example

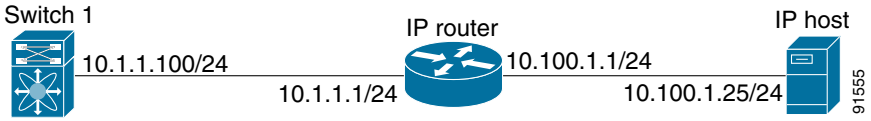

Note The port on the Ethernet switch to which the MDS Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in IOS is or the **set port host** in Catalyst OS. Refer to the configuration guide for your Ethernet switch.

To configure the Gigabit Ethernet interface for the example in Figure 45-1, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# ip address 10.1.1.100 255.255.255.0`	Enters the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

This section includes the following topics:

Configuring Interface Descriptions

See the "About Interface Descriptions" section on page 12-15 for details on configuring the switch port description for any interface.

Configuring Beacon Mode

See the "About Beacon Mode" section on page 12-17 for details on configuring the beacon mode for any interface.

Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

To configure autonegotiation, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# switchport auto-negotiate`	Enables autonegotiation for this Gigabit Ethernet interface (default).
	`switch(config-if)# no switchport auto-negotiate`	Disables autonegotiation for this Gigabit Ethernet interface.

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.

Note The minimum MTU size is 576 bytes.

Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

You do not need to explicitly issue the **shutdown** and **no shutdown** commands.

To configure the MTU frame size, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# switchport mtu 3000`	Changes the MTU size to 3000 bytes. The default is 1500 bytes.

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

To configure the promiscuous mode, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2` `switch(config-if)#`	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	`switch(config-if)# switchport promiscuous-mode on`	Enables promiscuous mode for this Gigabit Ethernet interface. The default is **off**.
	`switch(config-if)# switchport promiscuous-mode off`	Disables (default) promiscuous mode for this Gigabit Ethernet interface.
	`switch(config-if)# no switchport promiscuous-mode`	Disables (default) the promiscuous mode for this Gigabit Ethernet interface.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.

Note If the connection fails, verify the following, and ping the IP host again:
- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the `up` state.

Use the **ping** command to verify the Gigabit Ethernet connectivity (see Example 45-1). The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the "Using the ping and ping ipv6 Commands" section on page 2-15).

Use the **show interface gigabitethernet** command to verify if the Gigabit Ethernet interface is up.

Example 45-1 Verifying Gigabit Ethernet Connectivity

```
switch# ping 10.100.1.25
PING 10.100.1.25 (10.100.1.25): 56 data bytes
64 bytes from 10.100.1.25: icmp_seq=0 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=1 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=2 ttl=255 time=0.1 ms
--- 10.100.1.25 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.1 ms
```

VLANs

This section describes virtual LAN (VLAN) support in Cisco MDS SAN-OS and includes the following topics:

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

Note If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:
- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name (the `<slot-number>/<port-number>.<VLAN-ID>`).

Configuring the VLAN Subinterface

To configure a VLAN subinterface (VLAN ID), follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 2/2.100` `switch(config-if)#`	Specifies the subinterface on which 802.1Q is used (slot 2, port 2, VLAN ID 100). **Note** The subinterface number, 100 in this example, is the VLAN ID. The VLAN ID ranges from 1 to 4093.
Step 3	`switch(config-if)# ip address 10.1.1.101 255.255.255.0`	Enters the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see Table 45-1).

Table 45-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	

Note The configuration requirements in Table 45-1 also apply to Ethernet PortChannels.

Configuring Static IPv4 Routing

To configure static IPv4 routing (see Figure 45-1) through the Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1` `switch(config-if)#`	Enters the IP subnet (10.100.1.0 255.255.255.0) of the IP host and configures the next hop 10.1.1.1, which is the IPv4 address of the router connected to the Gigabit Ethernet interface.

Displaying the IPv4 Route Table

The **ip route interface** command takes the Gigabit Ethernet interface as a parameter and returns the route table for the interface. See Example 45-2.

Example 45-2 Displays the IP Route Table

```
switch# show ips ip route interface gig 8/1
Codes: C - connected, S - static
No default gateway
C 10.1.3.0/24 is directly connected, GigabitEthernet8/1
```

Connected (C) identifies the subnet in which the interface is configured (directly connected to the interface). Static (S) identifies the static routes that go through the router.

IPv4-ACLs

This section describes the guidelines for IPv4 access control lists (IPv4-ACLs) and how to apply them to Gigabit Ethernet interfaces.

This section includes the following topics:

- Gigabit Ethernet IPv4-ACL Guidelines, page 45-8
- Applying IPv4-ACLs on Gigabit Ethernet Interfaces, page 45-8

Note For information on creating IPv4-ACLs, see Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists."

Gigabit Ethernet IPv4-ACL Guidelines

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).

Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
  - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
  - The **established** option is ignored when you apply IPv4-ACLs containing this option to Gigabit Ethernet interfaces.
  - If an IPv4-ACL rule applies to a pre-existing TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B and an IPv4-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Tip If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. Chapter 34, “Configuring IPv4 and IPv6 Access Control Lists,” for information on configuring IPv4-ACLs.

Applying IPv4-ACLs on Gigabit Ethernet Interfaces

To apply an IPv4-ACL on a Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 3/1` `switch(config-if)#`	Configures a Gigabit Ethernet interface (3/1).
Step 3	`switch(config-if)# ip access-group SampleName`	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for both ingress and egress traffic (if the association does not exist already).
Step 4	`switch(config-if)# ip access-group SampleName1 in`	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for ingress traffic.
	`switch(config-if)# ip access-group SampleName2 out`	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for egress traffic (if the association does not exist already).

ARP Cache

Cisco MDS SAN-OS supports ARP cache for Gigabit Ethernet interface configured for IPv4. This section includes the following topics:

Displaying ARP Cache

You can display the ARP cache on Gigabit Ethernet interfaces.

Note Use the physical interface, not the subinterface, for all ARP cache commands.

Use the **show ips arp interface gigabitethernet** command to display the ARP cache on the Gigabit Ethernet interfaces. This command takes the Ethernet interface as a parameter and returns the ARP cache for that interface. See Example 45-3.

Example 45-3 Displays ARP Caches

```
switch# show ips arp interface gigabitethernet 7/1
Protocol         Address   Age (min)   Hardware Addr  Type   Interface
Internet        20.1.1.5         3     0005.3000.9db6  ARPA  GigabitEthernet7/1
Internet       20.1.1.10         7     0004.76eb.2ff5  ARPA  GigabitEthernet7/1
Internet       20.1.1.11        16     0003.47ad.21c4  ARPA  GigabitEthernet7/1
Internet       20.1.1.12         6     0003.4723.c4a6  ARPA  GigabitEthernet7/1
Internet       20.1.1.13        13     0004.76f0.ef81  ARPA  GigabitEthernet7/1
Internet       20.1.1.14         0     0004.76e0.2f68  ARPA  GigabitEthernet7/1
Internet       20.1.1.15         6     0003.47b2.494b  ARPA  GigabitEthernet7/1
Internet       20.1.1.17         2     0003.479a.b7a3  ARPA  GigabitEthernet7/1
...
```

Clearing ARP Cache

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache.

Use the **clear ips arp** command to clear the ARP cache. See Example 45-4 and Example 45-5.

Example 45-4 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 45-5 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```

Displaying IPv4 Statistics

Use the **show ips stats ip interface gigabitethernet** to display and verify IP v4 statistics. This command takes the main Ethernet interface as a parameter and returns the IPv4 statistics for that interface. See Example 45-6.

Note Use the physical interface, not the subinterface, to displayIPv4 statistics.

Example 45-6 Displays IPv4 Statistics

```
switch# show ips stats ip interface gigabitethernet 4/1
Internet Protocol Statistics for port GigabitEthernet4/1
    168 total received, 168 good, 0 error
    0 reassembly required, 0 reassembled ok, 0 dropped after timeout
    371 packets sent, 0 outgoing dropped, 0 dropped no route
    0 fragments created, 0 cannot fragment
```

Default Settings

Table 45-2 lists the default settings for IPv4 parameters.

Table 45-2 Default IPv4 Parameters

Parameters	Default
IPv4 MTU frame size	1500 bytes for all Ethernet ports.
Autonegotiation	Enabled.
Promiscuous mode	Disabled.

CHAPTER 46

Configuring IPv6 for Gigabit Ethernet Interfaces

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS SAN-OS. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.

This chapter includes the following sections:

Note For Cisco SAN-OS features that use IP addressing, refer to the chapters in this guide that describe those features for information on IPv6 addressing support.

Note To configure IP version 4 (IPv4) on a Gigabit Ethernet interface, see Chapter 45, "Configuring IPv4 for Gigabit Ethernet Interfaces".

About IPv6

IPv6 provides the following enhancements over IPv4:

- Allows networks to scale and provide global reachability.
- Reduces the need for private address and network address translation (NAT).
- Provides simpler autoconfiguration of addresses.

This section describes the IPv6 features supported by Cisco MDS SAN-OS and includes the following topics:

Extended IPv6 Address Space for Unique Addresses

IPv6 extends the address space by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides many more globally unique IP addresses. By being globally unique, IPv6 addresses enable global reachability and end-to-end security for networked devices, functionality that is crucial to the applications and services that are driving the demand for more addresses.

IPv6 Address Formats

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format x:x:x:x:x:x:x:x. The following are examples of IPv6 addresses:

2001:0DB8:7654:3210:FEDC:BA98:7654:3210

2001:0DB8:0:0:8:800:200C:417A

It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (::) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). Table 46-1 lists compressed IPv6 address formats.

Note Two colons (::) can be used only once in an IPv6 address to represent the longest successive hexadecimal fields of zeros.

Note The hexadecimal letters in IPv6 addresses are not case-sensitive.

Table 46-1 Compressed IPv6 Address Formats

IPv6 Address Type	Uncompressed Format	Compressed Format
Unicast	2001:0DB8:800:200C:0:0:0:417A	2001:0DB8:800:200C::417A
Multicast	FF01:0:0:0:0:0:0:101	FF01::101

IPv6 Address Prefix Format

An IPv6 address prefix, in the format *ipv6-prefix/prefix-length*, can be used to represent bit-wise contiguous blocks of the entire address space. The *ipv6-prefix* is specified in hexadecimal using 16-bit values between the colons. The *prefix-length* is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Address Type: Unicast

An IPv6 unicast address is an identifier for a single interface on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address. The Cisco MDS SAN-OS supports the following IPv6 unicast address types:

- Global addresses
- Link-local addresses

Global Addresses

Global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Figure 46-1 shows the structure of a global address.

Figure 46-1 Global Address Format

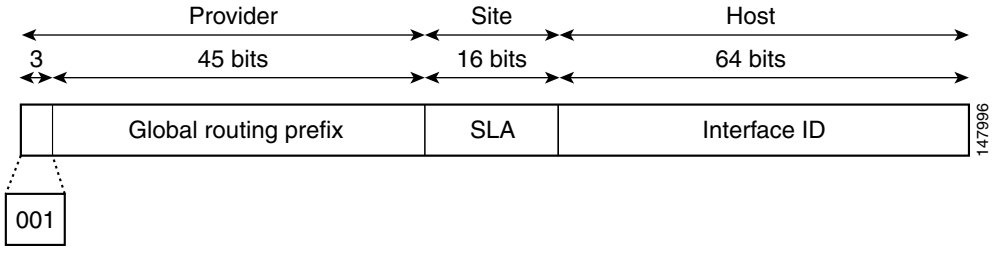

Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.

The aggregatable global address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA).The IETF decided to remove the TLS and NLA fields from the RFCs because these fields are policy-based. Some existing IPv6 networks deployed before the change might still be using networks based on the older architecture.

A 16-bit subnet field called the subnet ID could be used by individual organizations to create their own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

An interface ID is used to identify interfaces on a link. The interface ID must be unique to the link. They may also be unique over a broader scope. In many cases, an interface ID will be the same as, or based on, the link-layer address of an interface, which results in a globally unique interface ID. Interface IDs used in aggregatable global unicast and other IPv6 address types must be 64 bits long and constructed in the modified EUI-64 format.

Cisco MDS SAN-OS supports IEEE 802 interface types (for example, Gigabit Ethernet interfaces). The first three octets (24 bits) are taken from the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (MAC address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are taken from the last three octets of the MAC address. The construction of the interface ID is completed by setting the Universal/Local (U/L) bit—the seventh bit of the first octet—to a value of 0 or 1. A value of 0 indicates a locally administered identifier; a value of 1 indicates a globally unique IPv6 interface identifier (see Figure 46-2).

Figure 46-2 Interface Identifier Format

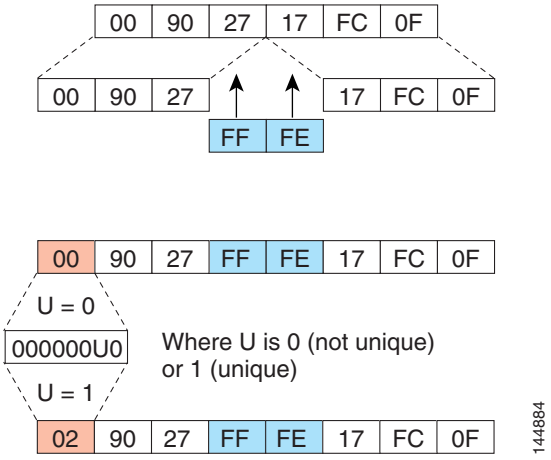

Link-Local Address

A link-local address is an IPv6 unicast address that is automatically configured on an interface using the link-local prefix FE80::/10 and the interface identifier in the modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate. Figure 46-3 shows the structure of a link-local address.

Figure 46-3 Link-Local Address Format

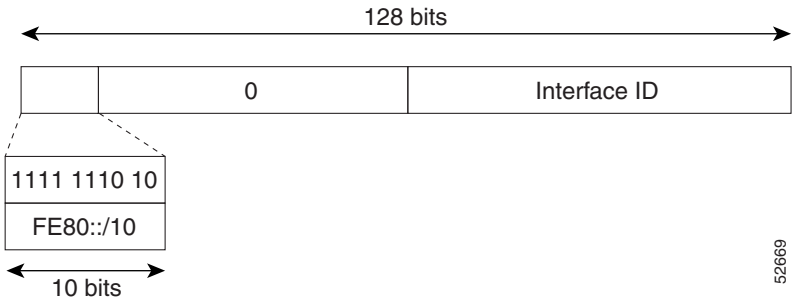

IPv6 Address Type: Multicast

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure 46-4 shows the format of the IPv6 multicast address.

Figure 46-4 IPv6 Multicast Address Format

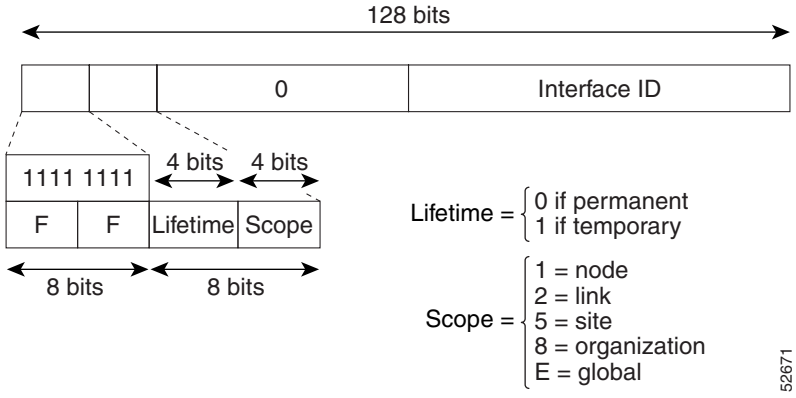

IPv6 hosts are required to join (receive packets destined for) the following multicast groups:

- All-node multicast group FF02::1.
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 concatenated with the low-order 24 bit of the unicast address.

  The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6

unicast address. (See Figure 46-5.) For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages.

Figure 46-5 IPv6 Solicited-Node Multicast Address Format

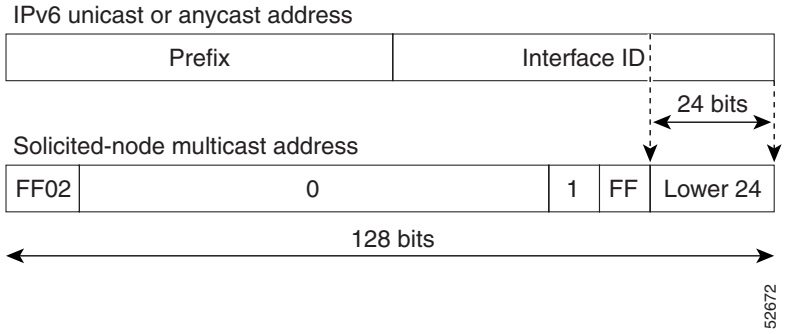

Note There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.

ICMP for IPv6

Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4—ICMP generates error messages such as ICMP destination unreachable messages, and informational messages such as ICMP echo request and reply messages. Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery process, path MTU discovery, and the Multicast Listener Discovery (MLD) protocol for IPv6. MLD is based on version 2 of the Internet Group Management Protocol (IGMP) for IPv4.

A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet. ICMP packets in IPv6 resemble a transport-layer packet in the sense that the ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet. Within IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is derived (computed by the sender and checked by the receiver) from the fields in the IPv6 ICMP packet and the IPv6 pseudoheader. The ICMPv6 Data field contains error or diagnostic information relevant to IP packet processing. Figure 46-6 shows the IPv6 ICMP packet header format.

Figure 46-6 IPv6 ICMP Packet Header Format

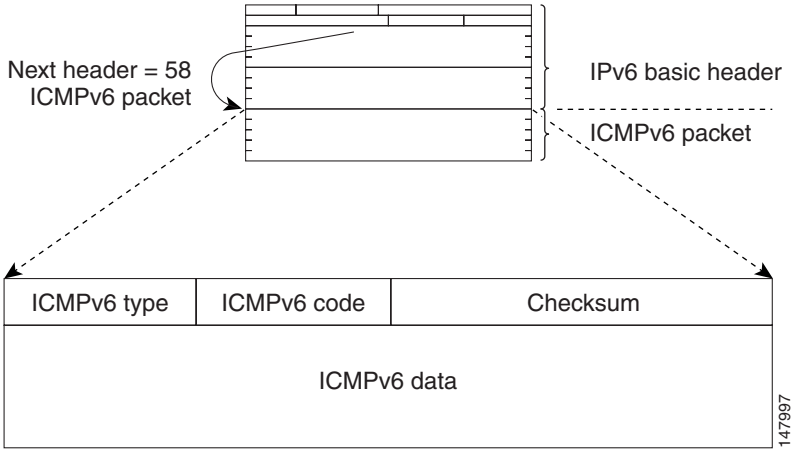

Path MTU Discovery for IPv6

As in IPv4, path MTU discovery in IPv6 allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, however, fragmentation is handled by the source of a packet when the path MTU of one link along a given data path is not large enough to accommodate the size of the packets. Having IPv6 hosts handle packet fragmentation saves IPv6 router processing resources and helps IPv6 networks run more efficiently.

Note In IPv4, the minimum link MTU is 68 octets, which means that the MTU size of every link along a given data path must support an MTU size of at least 68 octets.

In IPv6, the minimum link MTU is 1280 octets. We recommend using an maximum transmission unit (MTU) value of 1500 octets for IPv6 links.

IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers.

IPv6 Neighbor Solicitation and Advertisement Messages

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. (See Figure 46-7.) When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 46-7 IPv6 Neighbor Discovery—Neighbor Solicitation Message

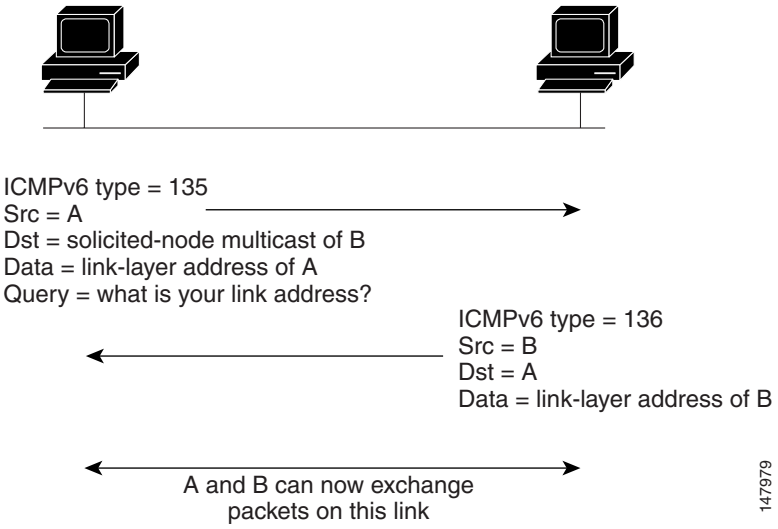

After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-node multicast address.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor, and is used for all paths between hosts and neighboring nodes (hosts or routers). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when the neighbor returns a positive acknowledgment indicating that it has received and processed packets previously sent to it. A positive acknowledgment could be from an upper-layer protocol such as TCP indicating that a connection is making forward progress (reaching its destination) or the receipt of a neighbor advertisement message in response to a neighbor solicitation message. If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Therefore, forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop router is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working. The return of a solicited neighbor advertisement message from the neighbor is a positive

acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.

Note A neighbor advertisement message that has the solicited flag set to a value of 0 must not be considered as a positive acknowledgment that the forward path is still working.

Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). Specifically, a node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

Every IPv6 unicast address (global or link-local) must be checked for uniqueness on the link; however, until the uniqueness of the link-local address is verified, duplicate address detection is not performed on any other IPv6 addresses associated with the link-local address.

Router Discovery

Router discovery performs both router solicitation and router advertisement. Router solicitations are sent by hosts to all-routers multicast addresses. Router advertisements are sent by routers in response to solicitations or unsolicited and contain default router information as well as additional parameters such as the MTU and hop limit.

IPv6 Stateless Autoconfiguration

All interfaces on IPv6 nodes must have a link-local address, which is automatically configured from the identifier for an interface and the link-local prefix FE80::/10. A link-local address enables a node to communicate with other nodes on the link and can be used to further configure the node.

Nodes can connect to a network and automatically generate site-local and global IPv6 address without the need for manual configuration or help of a server, such as a DHCP server. With IPv6, a router on the link advertises in router advertisement (RA) messages any site-local and global prefixes, and its willingness to function as a default router for the link. RA messages are sent periodically and in response to router solicitation messages, which are sent by hosts at system startup. (See Figure 46-8.)

Figure 46-8 IPv6 Stateless Autoconfiguration

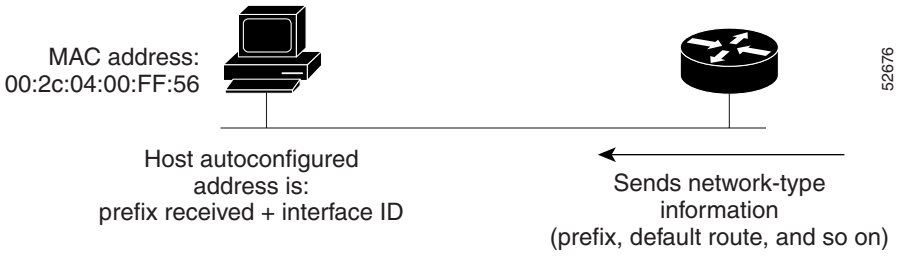

A node on the link can automatically configure site-local and global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure their uniqueness on the link. If the prefixes advertised in the RA messages are globally unique, then the IPv6 addresses configured by the node are also guaranteed to be globally unique. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message.

Dual IPv4 and IPv6 Protocol Stacks

The dual IPv4 and IPv6 protocol stack technique is one technique for a transition to IPv6. It enables gradual, one-by-one upgrades to applications running on nodes. Applications running on nodes are upgraded to make use of the IPv6 protocol stack. Applications that are not upgraded—they support only the IPv4 protocol stack—can coexist with upgraded applications on the same node. New and upgraded applications simply make use of both the IPv4 and IPv6 protocol stacks. (See Figure 46-9.)

Figure 46-9 Dual IPv4 and IPv6 Protocol Stack Technique

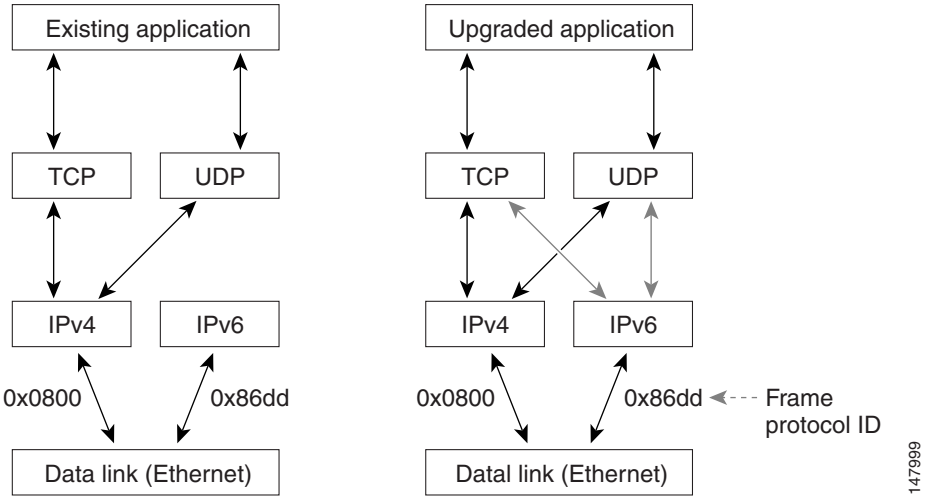

A new API has been defined to support both IPv4 and IPv6 addresses and DNS requests. An application can be upgraded to the new API and still use only the IPv4 protocol stack. The Cisco MDS SAN-OS supports the dual IPv4 and IPv6 protocol stack technique. When an interface is configured with both an IPv4 and an IPv6 address, the interface will accept and process both IPv4 and IPv6 traffic.

In Figure 46-10, an application that supports dual IPv4 and IPv6 protocol stacks requests all available addresses for the destination host name www.a.com from a DNS server. The DNS server replies with all available addresses (both IPv4 and IPv6 addresses) for www.a.com. The application chooses an address—in most cases, IPv6 addresses are the default choice—and connects the source node to the destination using the IPv6 protocol stack.

Figure 46-10 Dual IPv4 and IPv6 Protocol Stack Applications

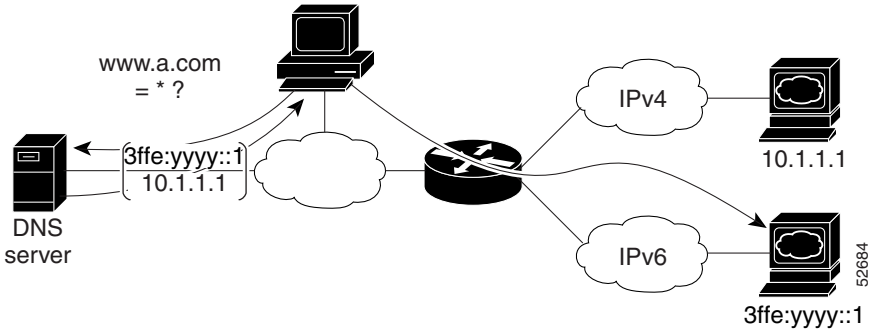

Configuring Basic Connectivity for IPv6

The tasks in this section explain how to implement IPv6 basic connectivity. Each task in the list is identified as either required or optional. This section includes the following topics:

- Configuring IPv6 Addressing and Enabling IPv6 Routing, page 46-11
- Configuring IPv4 and IPv6 Protocol Addresses, page 46-13
- Verifying Basic IPv6 Connectivity Configuration and Operation, page 46-13
- Clearing IPv6 Neighbor Discovery Cache, page 46-15

Configuring IPv6 Addressing and Enabling IPv6 Routing

This task explains how to assign IPv6 addresses to individual router interfaces and enable the processing of IPv6 traffic. By default, IPv6 addresses are not configured and IPv6 processing is disabled.

You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN

Note The *ipv6-address* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The *ipv6-prefix* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The *prefix-length* argument in the **ipv6 address** command is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:0:1:FF00::/104 for each unicast address assigned to the interface
- All-node link-local multicast group FF02::1

Note The solicited-node multicast address is used in the neighbor discovery process.

Note The maximum number of IPv6 addresses (static and autoconfigured) allowed on an interface is eight, except on the management (mgmt 0) interface where only one static IPv6 address can be configured.

To configure an IPv6 address on an interface and enable IPv6 routing, follow these steps:

	Command or Action	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 1/1` `switch(config-if)#`	Specifies a Gigabit Ethernet interface and enters interface configuration submode.
	`switch(config)# interface mgmt 0` `switch(config-if)#`	Specifies the management interface and enters interface configuration submode.
	`switch(config)# interface gigabitethernet 2/2.100` `switch(config-if)#`	Specifies a Gigabit Ethernet subinterface (VLAN ID) and enters interface configuration submode.
	`switch(config)# interface vsan 10` `switch(config-if)#`	Specifies a VSAN interface and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 address 2001:0DB8:800:200C::417A/64`	Assigns a unicast IPv6 address to the interface, automatically configures an IPv6 link-local address on the interface, and enables IPv6 processing on the interface.
	`switch(config-if)# ipv6 address autoconfig`	Enables autoconfiguration of the IPv6 link-local and unicast addresses on the interface, and enables IPv6 processing on the interface.
	`switch(config-if)# ipv6 enable`	Automatically configures an IPv6 link-local address on the interface while also enabling the interface for IPv6 processing. The link-local address can be used only to communicate with nodes on the same link.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.

	Command or Action	Purpose
Step 5	`switch(config-if)# exit` `switch(config)`	Exits interface configuration submode and returns to configuration mode.
Step 6	`switch(config)# ipv6 routing`	Enables the processing of IPv6 unicast datagrams.

Configuring IPv4 and IPv6 Protocol Addresses

When an interface in a Cisco networking device is configured with both an IPv4 and an IPv6 address, the interface can send and receive data on both IPv4 and IPv6 networks.

To configure an interface in a Cisco networking device to support both the IPv4 and IPv6 protocol stacks, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 1/1` `switch(config-if)#`	Specifies the interface, and enters interface configuration submode.
Step 3	`switch(config-if)# ip address 192.168.99.1 255.255.255.0`	Specifies a primary or secondary IPv4 address for an interface.
Step 4	`switch(config-if)# ipv6 address 2001:0DB8:c18:1::3/64`	Specifies the IPv6 network assigned to the interface and enables IPv6 processing on the interface. **Note** See the "Configuring IPv6 Addressing and Enabling IPv6 Routing" section for more information on configuring IPv6 addresses.
Step 5	`switch(config-if)# no shutdown`	Enables the interface.
Step 6	`switch(config-if)# exit` `switch(config)`	Exits interface configuration submode, and returns to configuration mode.
Step 7	`switch(config)# ipv6 routing`	Enables the processing of IPv6 unicast datagrams.

Verifying Basic IPv6 Connectivity Configuration and Operation

You can display information to verify the configuration and operation of basic IPv6 connectivity.

This section provides the following **show ipv6** command output examples:

- Example Output for the show ipv6 interface Command
- Example Output for the show ipv6 neighbours Command
- Example Output for the show ipv6 traffic Command

Example Output for the show ipv6 interface Command

In the following example, the **show ipv6 interface** command is used to verify that IPv6 addresses are configured correctly for the Gigabit Ethernet 6/1 interface.

```
switch# show ipv6 interface mgmt 0
mgmt0 is up
  IPv6 is enabled
  Global address(es):
    2172:22::180/64
  Link-local address(es):
    fe80::b8db:adff:feba:d074
  ND DAD is disabled
  ND reachable time is 30000 milliseconds
  ND retransmission time is 1000 milliseconds
  Stateless autoconfig for addresses disabled
  MTU is 1500 bytes
```

Example Output for the show ipv6 neighbours Command

In the following example, the **show ipv6 neighbours** command displays IPv6 neighbor discovery cache information for all interfaces.

```
switch# show ipv6 neighbours
R - Reachable, I - Incomplete, S - Stale, F - Failed, P - Probe, D - Delay
IPv6 Address                              Age  State Link-layer Addr     Interface
fe80::211:5dff:fe53:500a                  0    S     0011.5d53.500a      GigE6/1
fe80::211:5dff:fe53:500a                  0    S     0011.5d53.500a      GigE6/2
5000:1::250                               0    S     0011.5d53.500a      po 4
fe80::211:5dff:fe53:500a                  0    S     0011.5d53.500a      po 4
fe80::211:5dff:fe53:500a                  0    S     0011.5d53.500a      po 4
fe80::2d0:3ff:fe61:4800                   184  S     00d0.0361.4800      mgmt0
```

In the following example, the **show ipv6 neighbours interface** command displays IPv6 neighbor discovery cache information for the Gigabit Ethernet 6/1 interface.

```
switch# show ipv6 neighbours interface gigabitethernet 6/1
R - Reachable, I - Incomplete, S - Stale, F - Failed, P - Probe, D - Delay
IPv6 Address                              Age  State Link-layer Addr     Interface
fe80::211:5dff:fe53:500a                  0    S     0011.5d53.500a      GigE6/1
```

Example Output for the show ipv6 traffic Command

The **show ipv6 traffic** command displays IPv6 and ICMP statistics.

```
switch# show ipv6 traffic
IPv6 Statistics:
  Rcvd:  100 total, 0 local destination
         0 errors, 0 truncated, 0 too big
         0 unknown protocol, 0 dropped
         0 fragments, 0 reassembled
         0 couldn't reassemble, 0 reassembly timeouts
  Sent:  0 generated, 0 forwarded 0 dropped
         0 fragmented, 0 fragments created, 0 couldn't fragment

ICMPv6 Statistics:
  Rcvd:  100 total, 0 errors, 0 unreachables, 0 time exceeded
         0 too big, 0 param probs, 0 admin prohibs
         0 echos, 0 echo reply, 0 redirects
         0 group query, 0 group report, 0 group reduce
         0 router solicit, 69 router advert
         0 neighbor solicit, 31 neighbor advert
  Sent:  55 total, 0 errors, 0 unreachables, 0 time exceeded
         0 too big, 0 param probs, 0 admin prohibs
         0 echos, 0 echo reply, 0 redirects
```

```
0 group query, 20 group report, 2 group reduce
0 router solicit, 0 router advert
0 neighbor solicit, 33 neighbor advert
```

Clearing IPv6 Neighbor Discovery Cache

You can clear the IPv6 neighbor discovery cache using the **clear ipv6 neighbor** command in EXEC mode.

```
switch# clear ipv6 neighbor
```

Configuring Neighbor Discovery Parameters

You can configure the following neighbor discovery parameters:

- Duplicate address detection attempts
- Reachability time
- Retransmission timer

Note We recommend that you use the factory-defined defaults for these parameters.

This section includes the following topics:

Duplicate Address Detection Attempts

To configure the number of duplicate address detection attempts, follow these steps:

	Command or Action	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 3/1` `switch(config-if)#`	Specifies an interface and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 nd dad attempts 3`	Sets the duplicate address detection attempts count to 100. The range is 0 to 15.
Step 4	`switch(config-if)# no ipv6 nd dad attempts`	Reverts to the default value (0). **Note** When the attempt count is set to 0, neighbor discovery is disabled.

Reachability Time

To configure the reachability time, follow these steps:

	Command or Action	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 3/1` `switch(config-if)#`	Specifies an interface and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 nd reachability-time 10000`	Sets the reachability time to 10000 milliseconds. The range is 1000 to 3600000 milliseconds.
Step 4	`switch(config-if)# no ipv6 nd reachability-time`	Reverts to the default value (30000 milliseconds).

Retransmission Time

To configure the retransmission time, follow these steps:

	Command or Action	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface gigabitethernet 3/1` `switch(config-if)#`	Specifies an interface and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 nd retransmission-timer 20000`	Sets the retransmission time to 20000 milliseconds. The range is 1000 to 3600000 milliseconds.
Step 4	`switch(config-if)# no ipv6 nd retransmission-timer`	Reverts to the default value (1000 milliseconds).

Verifying Neighbor Discovery Parameter Configuration

The **show ipv6 interface** command displays the configuration of the neighbor discovery parameters.

```
switch# show ipv6 interface mgmt 0
mgmt0 is up
  IPv6 is enabled
  Global address(es):
    2003::1/64
  Link-local address(es):
    fe80::205:30ff:fe00:533e
  ND DAD is enabled, number of DAD attempts: 5
  ND reachable time is 50000 milliseconds
  ND retransmission time is 3000 milliseconds
  Stateless autoconfig for addresses disabled
```

Configuring IPv6 Static Routes

Cisco MDS SAN-OS supports static routes for IPv6. This section includes the following topics:

Configuring a IPv6 Static Route

You must manually configure IPv6 static routes and define an explicit path between two networking devices. IPv6 static routes are not automatically updated and must be manually reconfigured if the network topology changes.

To configure a IPv6 static route, follow these steps:

	Command or Action	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ipv6 route ::/0 gigabitethernet 3/1`	Configures a static default IPv6 route on a Gigabit Ethernet interface.
Step 3	`switch(config)# ipv6 route 2001:0DB8::/32 gigabitethernet 3/2`	Configures a fully specified IPv6 static route on a Gigabit Ethernet interface.

Verifying IPv6 Static Route Configuration and Operation

The **show ipv6 route** command displays the IPv6 route table for the switch.

```
switch# show ipv6 route

IPv6 Routing Table
Codes: C - Connected, L - Local, S - Static G - Gateway
G    ::/0
     via fe80::211:5dff:fe53:500a, GigabitEthernet6/1, distance 2
G    ::/0
     via fe80::2d0:3ff:fe61:4800, mgmt0, distance 2
C    2000::/64
     via ::, mgmt0
C    2172:22::/64
     via ::, mgmt0, distance 2
C    3000:3::/64
     via fe80::205:30ff:fe01:7ed6, GigabitEthernet4/1
C    3000:4::/64
     via fe80::205:30ff:fe01:7ed6, GigabitEthernet4/1.250
C    3000:5::/64
     via fe80::213:1aff:fee5:e69b, GigabitEthernet5/4
C    3000:6::/64
     via fe80::213:1aff:fee5:e69b, GigabitEthernet5/4.250
C    3000:7::/64
     via fe80::205:30ff:fe01:7ed7, GigabitEthernet4/2
C    3000:8::/64
     via fe80::205:30ff:fe01:7ed7, GigabitEthernet4/2.250
C    3000:9::/64
     via fe80::213:1aff:fee5:e69e, port-channel 3
C    3000:10::/64
     via fe80::213:1aff:fee5:e69e, port-channel 3.250
C    5000:1::/64
     via fe80::205:30ff:fe01:3917, GigabitEthernet6/2
C    5000:1::/64
     via fe80::205:30ff:fe01:3918, port-channel 4
C    6000:1:1:1::/64
```

```
     via fe80::205:30ff:fe01:3916, GigabitEthernet6/1
C    7000:1::/64
     via fe80::205:30ff:fe01:3917, GigabitEthernet6/2.250
C    7000:1::/64
     via fe80::205:30ff:fe01:3918, port-channel 4.250
C    7000:1:1:1::/64
     via fe80::205:30ff:fe01:3917, GigabitEthernet6/2, distance 2
L    fe80::/10
     via ::
L    ff00::/8
     via ::
```

Gigabit Ethernet IPv6-ACL Guidelines

Tip If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists" for information on configuring IPv6-ACLs.

Follow these guidelines when configuring IPv6-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).

  Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv6-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
  - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
  - The **established** option is ignored when you apply IPv6-ACLs containing this option to Gigabit Ethernet interfaces.
  - If an IPv6-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example, if there is an existing TCP connection between A and B and an IPv6-ACL that specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

See Chapter 34, "Configuring IPv4 and IPv6 Access Control Lists" for information on applying IPv6-ACLs to an interface.

Transitioning from IPv4 to IPv6

Cisco MDS SAN-OS does not support any transitioning mechanisms from IPv4 to IPv6. However, you can use the transitioning schemes in the Cisco router products for this purpose. For information on configuring Cisco routers to transition your network, refer to the *Implementing Tunneling for IPv6* document in the *Cisco IOS IPv6 Configuration Library*.

Displaying IPv6 Information

Use the **show ips ipv6 neighbours interface** command for information about IPv6 neighbors for an interface.

```
switch# show ips ipv6 neighbours interface gigabitethernet 6/1
IPv6 Address                          Age (min)  Link-layer Addr  State  Interface
fe80::211:5dff:fe53:500a              0          0011.5d53.500a      S   Gigabi tEthernet6/1
```

Use the **show ips ipv6 prefix-list interface** command for information about IPv6 prefixes for an interface.

```
switch# show ips ipv6 prefix-list interface gigabitethernet 6/1
Prefix                               Prefix-len   Addr
     Valid   Preferred
6000:1:1:1::                          64         ::
       2592000    604800
```

Use the **show ips ipv6 interface** command for information about the IPv6 routes for an interface.

```
switch# show ips ipv6 route interface gigabitethernet 6/1
IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, G - Gateway, M - Multicast
C 6000:1:1:1::/64 is directly connected, GigabitEthernet6/1
C 6000:1:1:1::/64 is directly connected, GigabitEthernet6/1
C fe80::/64 is directly connected, GigabitEthernet6/1
M ff02::/32 is multicast, GigabitEthernet6/1
G ::/0 via fe80::211:5dff:fe53:500a, GigabitEthernet6/1
```

Use the **show ips ipv6 routers interface** command for information about IPv6 routers for an interface.

```
switch# show ips ipv6 routers interface gigabitethernet 6/1
Addr                              Lifetime   Expire
fe80::211:5dff:fe53:500a          1800      1781
```

Use the **show ips ipv6 traffic interface** command for information about IPv6 traffic statistics for an interface.

```
switch# show ips ipv6 traffic interface gigabitethernet 6/1
IPv6 statistics:
  Rcvd:  5094 total
         0 bad header, 0 unknown option, 0 unknown protocol
         0 fragments, 0 total reassembled
         0 reassembly timeouts, 0 reassembly failures
  Sent:  13625 generated
         0 fragmented into 0 fragments, 0 failed
         2 no route
ICMP statistics:
  Rcvd: 1264 input, 0 checksum errors, 0 too short
         0 unknown info type, 0 unknown error type
         unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
         parameter: 0 error, 0 header, 0 option
         0 hopcount expired, 0 reassembly timeout,0 too big
         0 echo request, 0 echo reply
         734 group query, 0 group report, 0 group reduce
         0 router solicit, 528 router advert, 0 redirects
         0 neighbor solicit, 2 neighbor advert
  Sent: 6045 output, 0 rate-limited
         unreach: 0 routing, 0 admin, 0 neighbor, 1160 address, 0 port
         parameter: 0 error, 0 header, 0 option
         0 hopcount expired, 0 reassembly timeout,0 too big
         0 echo request, 0 echo reply
         0 group query, 1466 group report, 0 group reduce
```

```
1 router solicit, 0 router advert, 0 redirects
3412 neighbor solicit, 6 neighbor advert
```

Default Settings

Table 46-2 lists the default settings for IPv6 parameters.

Table 46-2 Default IPv6 Parameters

Parameters	Default
IPv6 processing	Disabled.
Duplicate address detection attempts	0 (neighbor discovery disabled).
Reachability time	1000 milliseconds.
Retransmission time	30000 milliseconds.
IPv6-ACLs	None.

PART 7

Intelligent Storage Services

CHAPTER 47

Configuring SCSI Flow Services and Statistics

Intelligent Storage Services are features supported on the Storage Services Module (SSM). The Storage Services Module (SSM) supports SCSI flow services and SCSI flow statistics. Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.0(2b) and later include the following topics:

SCSI Flow Services

A SCSI initiator/target combination is a SCSI flow. SCSI flow services provide enhanced features for SCSI flows, such as write acceleration and flow monitoring for statistics gathering on an SSM.

This section includes the following topics:

About SCSI Flow Services

A SCSI initiator/target combination is a SCSI flow. SCSI flow services provide enhanced features for SCSI flows, such as write acceleration and flow monitoring for statistics gathering on an SSM.

Functionally, the SCSI flow services functional architecture consists of the following components:

- SCSI flow manager (SFM) on the supervisor
- SCSI flow configuration CLI on the supervisor
- SCSI flow configuration client on the Control Path Processor (CPP) of an SSM
- SCSI flow feature set support on the Data Path Processor (DPP) of an SSM

Figure 47-1 shows an example of the SCSI flow services functional architecture.

Figure 47-1 SCSI Flow Services Functional Architecture

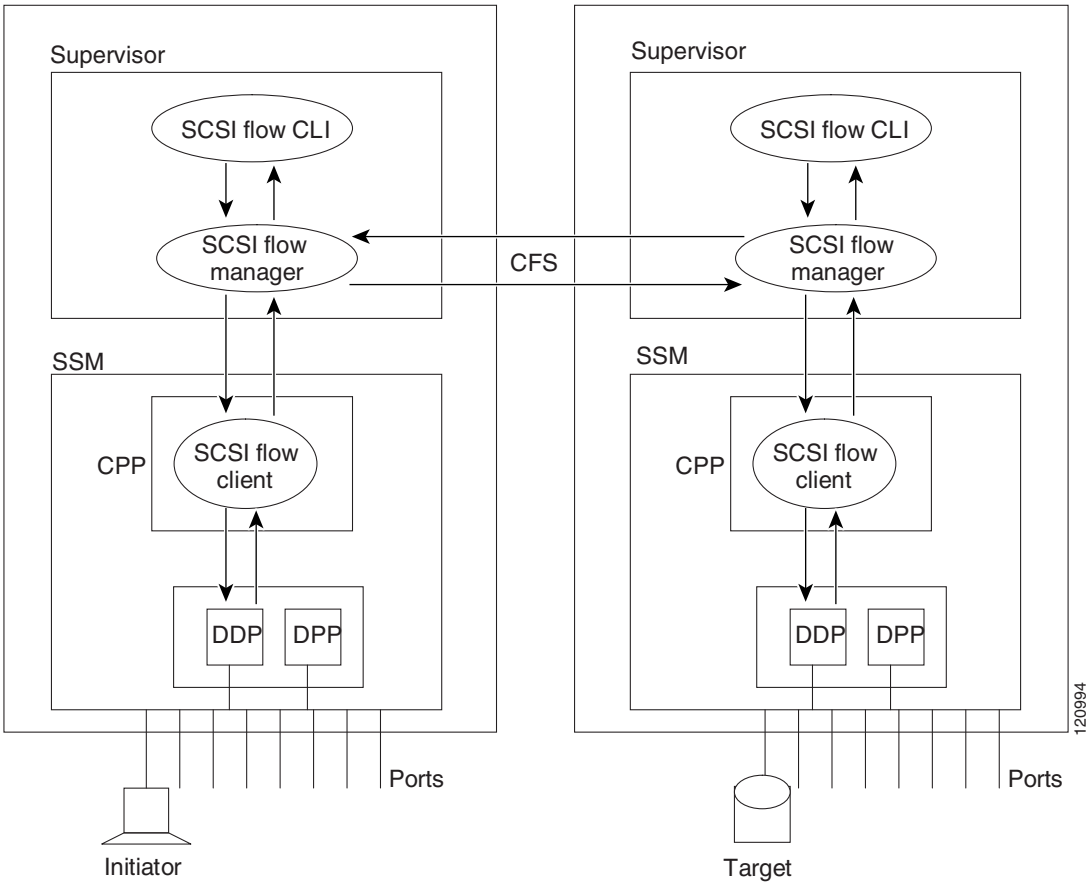

Note The SCSI target and initiator must be connected to different SSMs on different switches.

Note For statistics monitoring, the target device is not required to be connected to an SSM.

SCSI Flow Manager

The SCSI flow manager (SFM) resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events. The SFM registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SFM on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SFM to validate target parameters and program information on the target side.

SCSI Flow Configuration Client

A SCSI flow configuration client (SFCC) resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.

SCSI Flow Data Path Support

The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features such as Fibre Channel write acceleration and statistics monitoring.

Configuring SCSI Flow Services

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN
- Flow feature set consisting of Fibre Channel write acceleration and statistics monitoring.

The SCSI flow specification is a distributed configuration because the SCSI initiator and the target might be physically connected to SSMs on two different switches located across the fabric. The configuration does not require information to identify either the switch name or the SSM slot location for either the initiator or the target. The manual SCSI flow configuration is performed only at the initiator side. This simplifies the configuration process. The initiator switch sends the configuration to the SFM on the target switch using CFS. No SCSI flow configuration is necessary on the target switch.

Enabling SCSI Flow Services

You can enable SCSI flow services either on the entire SSM or on groups of four interfaces.

Enabling SCSI flow services on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

Note Fibre Channel write acceleration can only be provisioned on the entire SSM, not a group of interfaces on the SSM.

To enable SCSI flow services, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ssm enable feature scsi-flow module 2`	Enables SCSI flow services on the SSM in slot 2.
	`switch(config)# no ssm enable feature scsi-flow module 2`	Disables SCSI flow services on the SSM in slot 2. The default is disabled.
	`switch(config)# no ssm enable feature scsi-flow force module 2`	Forces the switch to disable SCSI flow services on the SSM in slot 2. The default is disabled.
Step 3	`switch(config)# ssm enable feature scsi-flow interface fc 2/5 - 8`	Enables SCSI flow services on interface 5 through 8 on the SSM in slot 2. **Note** Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29.
	`switch(config)# no ssm enable feature scsi-flow interface fc 2/5 - 8`	Disables SCSI flow services on interface 5 through 8 on the SSM in slot 2. The default is disabled.
	`switch(config)# no ssm enable feature scsi-flow force interface fc 2/5 - 8`	Forces the switch to disable SCSI flow services on the interface 5 through 8 on the SSM in slot 2.

Enabling SCSI Flow Configuration Distribution

To enable SCSI flow configuration distribution using CFS, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# scsi-flow distribute`	Enables SCSI flow configuration distribution through CFS. The default is enabled.
	`switch(config)# no scsi-flow distribute`	Disables CFS distribution for SCSI flow configuration.

Configuring SCSI Flow Identifiers

A SCSI flow identifier is unique on a switch and is chosen by the user, like VSAN identifiers. To configure a SCSI flow identifier, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# scsi-flow flow-id 3 initiator-vsan 2 initiator-pwwn 21:00:00:e0:8b:07:5f:aa target-vsan 4 target-pwwn 2a:20:00:05:30:00:77:e0`	Configures SCSI flow identifier 3 using the pWWNs of the initiator and the target. The flow identifier range is 1 to 65535.
	`switch(config)# no scsi-flow flow-id 3 initiator-vsan 2`	Removes SCSI flow identifier 3.

SCSI Flow Statistics

This section includes the following topics:

About SCSI Flow Statistics

The statistics that can be collected for SCSI flows include the following:

- SCSI reads
  - Number of I/Os
  - Number of I/O blocks
  - Maximum I/O blocks
  - Minimum I/O response time
  - Maximum I/O response time
- SCSI writes
  - Number of I/Os
  - Number of I/O blocks
  - Maximum I/O blocks
  - Minimum I/O response time
  - Maximum I/O response time
- Other SCSI commands (not read or write)
  - Test unit ready
  - Report LUN
  - Inquiry
  - Read capacity

  - Mode sense
  - Request sense
- Errors
  - Number of timeouts
  - Number of I/O failures
  - Number of various SCSI status events
  - Number of various SCSI sense key errors or events

To take advantage of this feature, only the initiator must be directly attached to an SSM.

Note The SCSI flow statistics feature requires the Enterprise Package license installed only on the initiator switches.

Note For SCSI flow statistics, the initiator must connect to an SSM on a Cisco MDS switch while the target can connect to any other switch in the fabric. The SCSI flow initiator and target cannot connect to the same switch.

Configuring SCSI Flow Statistics

This section includes the following topics:

Enabling SCSI Flow Statistics

To enable SCSI flow statistics monitoring, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# scsi-flow flow-id 3 statistics`	Enables statistics monitoring on SCSI flow identifier 3.
	`switch(config)# no scsi-flow flow-id 3 statistics`	Disables statistics monitoring on SCSI flow identifier 3. The default is disabled.

Clearing SCSI Flow Statistics

Use the c**lear device-name statistics flow-id** command to clear SCSI flow statistics (for debugging purposes):

```
switch# clear scsi-flow statistics flow-id 3
```

Displaying SCSI Flow Services Information

Use the **show scsi-flow** command to display information about SCSI flow services (see Example 47-1 to Example 47-5).

Example 47-1 Displays Applications Provisioned on an SSM

```
switch# show ssm provisioning
Module   Ports       Application        Provisioning Status
--------------------------------------------------------
   4      1-32           scsi-flow            success
```

Example 47-2 Displays SCSI Flow Services Configuration for All SCSI Flow Identifiers

```
switch# show scsi-flow
Flow Id: 3
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:7f:7d
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
          Statistics enabled
            Configuration Status:  success

Flow Id: 4
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:a7:89
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
```

Example 47-3 Displays SCSI Flow Services Configuration for a Specific SCSI Flow Identifier

```
switch# show scsi-flow flow-id 3
Flow Id: 3
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:7f:7d
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
          Statistics enabled
            Configuration Status:  success
```

Example 47-4 Displays SCSI Flow Services Statistics for All SCSI Flow Identifiers

```
switch# show scsi-flow statistics

Stats for flow-id 4 LUN=0x0000
------------------------------
 Read Stats
  I/O Total count=2
  I/O Timeout count=0
  I/O Total block count=4
  I/O Max block count=2
  I/O Min response time=5247 usec
  I/O Max response time=10160 usec
  I/O Active Count=0

 Write Stats
  I/O Total count=199935
  I/O Timeout count=0
  I/O Total block count=12795840
  I/O Max block count=64
  I/O Min response time=492 usec
  I/O Max response time=10056529 usec
  I/O Active Count=16

 Non Read-Write Stats
  Test Unit Ready=4
  Report LUN=38
  Inquiry=50
  Read Capacity=3
  Mode Sense=0
  Request Sense=0

 Total Stats
  Rx Frame Count=3792063
  Rx Frame Byte Count=6549984752
  Tx Frame Count=3792063
  Tx Frame Byte Count=6549984752
```

```
 Error Stats
  SCSI Status Busy=0
  SCSI Status Reservation Conflict=0
  SCSI Status Task Set Full=0
  SCSI Status ACA Active=0
  Sense Key Not Ready=0
  Sense Key Medium Error=0
  Sense Key Hardware Error=0
  Sense Key Illegal Request=0
  Sense Key Unit Attention=28
  Sense Key Data Protect=0
  Sense Key Blank Check=0
  Sense Key Copy Aborted=0
  Sense Key Aborted Command=0
  Sense Key Volume Overflow=0
  Sense Key Miscompare=0
```

Example 47-5 Displays SCSI Flow Services Statistics for a Specific SCSI Flow Identifier

```
switch# show scsi-flow statistics flow-id 4

Stats for flow-id 4 LUN=0x0000
------------------------------
 Read Stats
  I/O Total count=2
  I/O Timeout count=0
  I/O Total block count=4
  I/O Max block count=2
  I/O Min response time=5247 usec
  I/O Max response time=10160 usec
  I/O Active Count=0

 Write Stats
  I/O Total count=199935
  I/O Timeout count=0
  I/O Total block count=12795840
  I/O Max block count=64
  I/O Min response time=492 usec
  I/O Max response time=10056529 usec
  I/O Active Count=16

 Non Read-Write Stats
  Test Unit Ready=4
  Report LUN=38
  Inquiry=50
  Read Capacity=3
  Mode Sense=0
  Request Sense=0

 Total Stats
  Rx Frame Count=3792063
  Rx Frame Byte Count=6549984752
  Tx Frame Count=3792063
  Tx Frame Byte Count=6549984752

 Error Stats
  SCSI Status Busy=0
  SCSI Status Reservation Conflict=0
  SCSI Status Task Set Full=0
  SCSI Status ACA Active=0
  Sense Key Not Ready=0
  Sense Key Medium Error=0
  Sense Key Hardware Error=0
```

```
Sense Key Illegal Request=0
Sense Key Unit Attention=28
Sense Key Data Protect=0
Sense Key Blank Check=0
Sense Key Copy Aborted=0
Sense Key Aborted Command=0
Sense Key Volume Overflow=0
Sense Key Miscompare=0
```

Default Settings

Table 47-1 lists the default settings for SCSI flow services and SCSI flow statistics parameters.

Table 47-1 Default Intelligent Storage Services Parameters

Parameters	Default
SCSI flow services	Disabled.
SCSI flow services distribution	Enabled.
SCSI flow statistics	Disabled.

CHAPTER 48

Configuring Fibre Channel Write Acceleration

The Storage Services Module (SSM) supports Fibre Channel write acceleration on Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 2.0(2b) and later.

This chapter includes the following sections:

Fibre Channel Write Acceleration

Fibre Channel write acceleration minimizes application latency or reduces transactions per second over long distances. For synchronous data replication, Fibre Channel write acceleration increases the distance of replication or reduces effective latency to improve performance. To take advantage of this feature, both the initiator and target devices must be directly attached to an SSM.

This section includes the following topics:

About Fibre Channel Write Acceleration

The Fibre Channel write acceleration feature also allows the configuration of the buffer count. You can change the number of 2-KB buffers reserved on the target side DPP for a SCSI flow.

You can estimate the number of buffers to configure using the following formula:

(Number of concurrent SCSI writes * size of SCSI writes in bytes) / FCP data frame size in bytes

For example, HDS TrueCopy between HDS 9970s uses 1-KB FCP data frames. You perform an initial sync for a 16-LUN TrueCopy group with 15 tracks, or 768-KB per LUN, requires approximately 16*(768*1024)/1024 or 12248 write buffers.

Note The Fibre Channel write acceleration feature requires the Enterprise Package license installed on both the initiator and target switches.

Note The initiator and target cannot connect to the same Cisco MDS switch. Fibre Channel write acceleration requires that the initiator and target must each connect to an SSM module installed on different Cisco MDS switches.

Enabling Fibre Channel Write Acceleration

To enable Fibre Channel write acceleration, and optionally modify the number of write acceleration buffers, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ssm enable feature scsi-flow module 2`	Enables SCSI flow services on the SSM in slot 2. **Note** Fibre Channel write acceleration can only be configured on all interfaces on the SSM, not on groups of interfaces.
Step 3	`switch(config)# scsi-flow flow-id 3 initiator-vsan 2 initiator-pwwn 21:00:00:e0:8b:07:5f:aa target-vsan 4 target-pwwn 2a:20:00:05:30:00:77:e0`	Configures SCSI flow identifier 3 using the pWWNs of the initiator and the target. The flow identifier range is 1 to 65535.
Step 4	`switch(config)# scsi-flow distribute`	Enables CFS distribution for the SCSI flow. **Note** No CFS configuration commit operation is required for SCSI flow. The SCSI flow manager uses CFS for target validation.
Step 5	`switch(config)# scsi-flow flow-id 3 write-acceleration`	Enables Fibre Channel write acceleration for SCSI flow identifier 3.
	`switch(config)# no scsi-flow flow-id 3 write-acceleration`	Disables SCSI flow write acceleration for SCSI flow identifier 3. The default is disabled.
Step 6	`switch(config)# scsi-flow flow-id 3 write-acceleration buffer 2048`	Enables Fibre Channel write acceleration for SCSI flow identifier 3 and sets the number of buffers to 2048. The range is 1 to 40000.
	`switch(config)# no scsi-flow flow-id 3 write-acceleration buffer 1024`	Reverts to the default number of write acceleration buffers. The default is 1024.

Displaying Fibre Channel Write Acceleration Information

Use the **show scsi-flow** command to display information about the status of the Fibre Channel write acceleration configuration (see Example 48-1 and Example 48-2).

Example 48-1 Displays Fibre Channel Write Acceleration Configuration for All SCSI Flow Identifiers

```
switch# show scsi-flow
Flow Id: 3
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:7f:7d
```

```
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
          Statistics enabled
            Configuration Status:  success

Flow Id: 4
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:a7:89
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
          Statistics enabled
            Configuration Status:  success
```

Example 48-2 Displays Fibre Channel Write Acceleration Configuration for a Specific SCSI Flow Identifier

```
switch# show scsi-flow flow-id 3
Flow Id: 3
        Initiator VSAN: 101
        Initiator WWN: 21:00:00:e0:8b:05:76:28
        Target VSAN: 102
        Target WWN: 21:00:00:20:37:38:7f:7d
        Target LUN: ALL LUNs
        Flow Verification Status:
        -------------------------
           Initiator Verification Status:   success
           Target Verification Status:      success
           Initiator Linecard Status:       success
           Target Linecard Status:          success
        Feature Status:
        ---------------
          Write-Acceleration enabled
            Write-Acceleration Buffers: 1024
            Configuration Status:  success
          Statistics enabled
            Configuration Status:  success
```

Default Settings

Table 48-1 lists the default settings for Fibre Channel write acceleration parameters.

Table 48-1 Default Fibre Channel Write Acceleration Parameters

Parameters	Default
Fibre Channel write acceleration	Disabled.
Fibre Channel write acceleration buffers	1024.

CHAPTER 49

Configuring SANTap

The Storage Services Module (SSM) supports SANTap in Cisco MDS SAN-OS Release 2.0(2b) and later.

This chapter includes the following sections:

About SANTap

The SANTap feature allows third-party data storage applications, such as long distance replication and continuous backup, to be integrated into the SAN. The protocol-based interface that is offered by SANTap allows easy and rapid integration of the data storage service application because it delivers a loose connection between the application and an SSM, which reduces the effort needed to integrate applications with the core services being offered by the SSM. See Figure 49-1.

Figure 49-1 Integrating Third-Party Storage Applications in a SAN

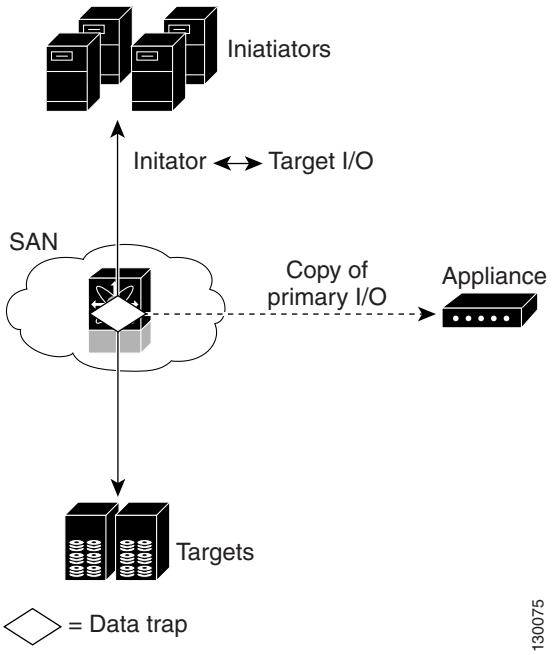

SANTap has a control path and a data path. The control path handles requests that create and manipulate replication sessions sent by an appliance. The control path is implemented using a SCSI-based protocol. An appliance sends requests to a Control Virtual Target (CVT), which the SANTap process creates and monitors. Responses are sent to the control LUN on the appliance. SANTap also allows LUN mapping to appliance virtual targets (AVTs). You can have a maximum of 512 target LUNs.

SANTap does not require reconfiguration of either the host or target when introducing SANTap-based applications. Also, neither the host initiator nor the target is required to be directly connected to an SSM. This is accomplished by assigning Cisco-specific WWNs to the virtual initiators (VIs) and Data Virtual Targets (DVTs). A host initiator or a target can be connected directly to an SSM. However, you must partition the SAN using VSANs.

You must configure the host initiator and the DVT in one VSAN and configure the VI and the target in another VSAN. See Figure 49-2.

Figure 49-2 SANTap Proxy Mode-2 Example

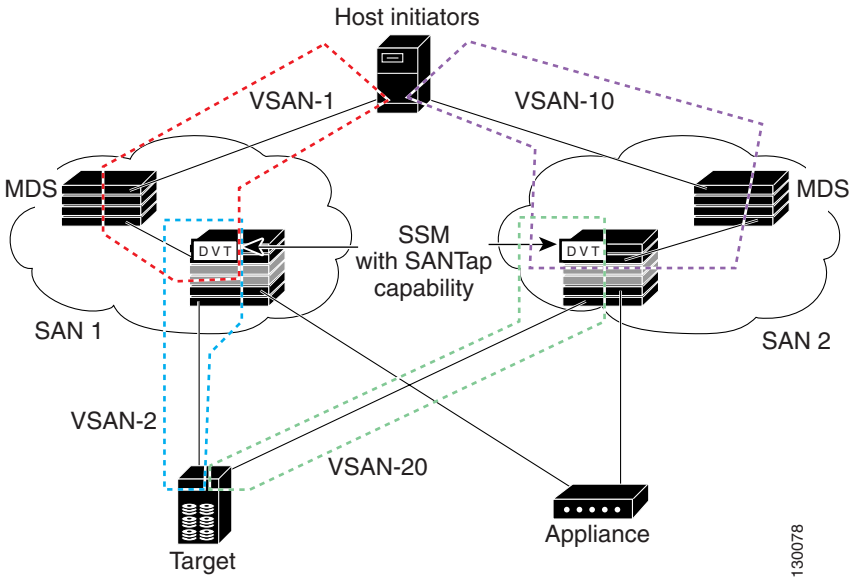

You can use SANTap to remove your appliance-based storage applications from the primary data path in your SAN. Removing these applications from the primary data path prevents them from compromising the security, availability, and performance of your SAN. SANTap copies the data at line speed and makes it available to other storage applications; these storage applications are prevented from affecting your SAN while maintaining the integrity of the data the storage applications need.

Dynamic LUN is a feature introduced in the Cisco SAN OS release 3.2(1). When one or more LUNs are removed or added on the backend target during the periodic scan, SANTap automatically uninstalls the deleted DVT LUNs and installs any additional LUNs. Uninstallation of the deleted DVT LUNs is done even if the total number of LUNs remains the same.

In previous releases, when the set of LUNs changed on the target, the original LUN list was displayed on the DVT. The new and changed LUNs were not reflected on the DVT. However, if the total number of LUNs increased, then the additional LUNs were installed and displayed on the host.

Prior to Cisco SAN OS release 3.2(1), a user had the following options for displaying the LUN list on DVT:

- Shut the host interface- Purge the DVT LUNs for the IT pair using CLI. All the LUNs for the existing IT pair were removed, and the correct set of LUNs was recreated when the host logs in.
- Reload the SSM- Works only if there are no sessions and AVT LUNs present.

64-Bit LUN Support- In Cisco SAN OS release 3.2(1) or later releases, SANTap supports 64-bit LUNs on the target.

The following CLI commands are used to obtain the mapping between the host-side LUN and the target-side LUN:

```
Switch# show santap module <num> dvtlun
Switch# show santap module <num> dvtlun brief
Module# show santap vttbl dvt <dvt_wwn> host <host_wwn>
```

Configuring SANTap

This section includes the following topics:

Enabling SANTap

SANTap can be enabled on an entire SSM or it can be enabled on a group of four ports on an SSM. Enabling SANTap on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

To enable the SANTap feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ssm enable feature santap module 4`	Enables the SANTap application on the entire SSM.
	`switch(config)# no ssm enable feature santap module 4`	Disables the SANTap application on the entire SSM in slot 4.
	`switch(config)# no ssm enable feature santap force module 4`	Forces the switch to disable the SANTap application on the entire SSM in slot 4.
Step 3	`switch(config)# ssm enable feature santap interface fc 4/1 - 4`	Enables the SANTap application on ports 1 through 4 on the SSM. **Note** Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29.
	`switch(config)# no ssm enable feature santap interface fc 4/1 - 4`	Disables the SANTap application on ports 1 through 4 on the SSM in slot 4.
	`switch(config)# no ssm enable feature santap force interface fc 4/1 - 4`	Forces the switch to disable the SANTap application on ports1 through 4 on the SSM in slot 4.
Step 4	`switch(config)# santap module 4 appl-vsan 10`	Enables SANTap on the SSM in slot 4 and on VSAN 10.
	`switch(config)# no santap module 4 appl-vsan 10`	Disables SANTap.

Note You cannot simultaneously configure the intelligent services SANTap and NASB on a single SSM.

Configuring DVTs

To configure a DVT, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# santap module 2 dvt target-pwwn 50:06:0e:80:03:81:32:36 target-vsan 9 dvt-name MYDVT dvt-vsan 12`	Configures the pWWN, target VSAN (which contains the target and VI), DVT name, and DVT VSAN (which contains the host and the CVT).
	`switch(config)# santap module 2 dvt target-pwwn 50:06:0e:80:03:81:32:36 target-vsan 9 dvt-name MYDVT dvt-vsan 12 dvt-port 1`	Configures the pWWN, target VSAN, DVT name, DVT VSAN, and DVT port. The DVT port maps to one of the ports on the SSM. You can assign a port for explicit load balancing or not assign a port, which allows the SSM select the port and handle the load balancing (default).
	`switch(config)# santap module 2 dvt target-pwwn 50:06:0e:80:03:81:32:36 target-vsan 9 dvt-name MYDVT dvt-vsan 12 lun-size-handling 1`	Configures the pWWN, target VSAN, DVT name, DVT VSAN, and LUN size handling flag (enabled). Enabling the LUN size handling flag allows special LUN resize handling by the vendor. The default LUN size handling flag value is 0 (disabled).
	`switch(config)# santap module 2 dvt target-pwwn 50:06:0e:80:03:81:32:36 target-vsan 9 dvt-name MYDVT dvt-vsan 12 io-timeout 20`	Configures the pWWN, target VSAN, DVT name, DVT VSAN, and IO timeout value in seconds. The IO timeout determines the interval after which to time out I/Os on the target side. The range is 10 to 200 seconds and the default value is 10 seconds.
	`switch(config)# no santap module 2 dvt target-pwwn 50:06:0e:80:03:81:32:36`	Removes the DVT configuration.

In Cisco SAN OS release 3.2(1) or later releases, SANTap supports 32 host initiators per DVT.

Displaying SANTap Information

Use the **show santap module** command to display information about SANTap (see Example 49-1 to Example 49-8).

Example 49-1 Displays SANTap CVT Information

```
switch# show santap module 2 cvt

CVT Information :
        cvt pwwn     = 23:4f:00:0d:ec:09:3c:02
        cvt nwwn     = 23:9d:00:0d:ec:09:3c:02
        cvt id       = 135895180
        cvt xmap_id  = 135895212
        cvt vsan     = 8
```

```
        cvt name     =
```

Example 49-2 Displays SANTap DVT Information

```
switch# show santap module 2 dvt

DVT Information :
       dvt pwwn     = 50:06:0e:80:03:81:32:36
       dvt nwwn     = 50:06:0e:80:03:81:32:36
       dvt id       = 136773180
       dvt mode     = 3
       dvt vsan     = 12
       dvt if_index = 0x1080000
       dvt fp_port  = 1
       dvt name     = MYDVT
       dvt tgt-vsan = 9
       dvt io timeout            = 10 secs
       dvt lun size handling     = 0
       dvt app iofail behaviour  = 1
       dvt quiesce behavior      = 1
       dvt tgt iofail behavior   = 0
       dvt appio failover time   = 50 secs
       dvt inq data behavior     = 0
```

Example 49-3 Displays SANTap DVT LUN Information

```
switch# show santap module 2 dvtlun

 DVT LUN Information :
       dvt pwwn     = 22:00:00:20:37:88:20:ef
       dvt lun      = 0x0
       xmap id      = 8
       dvt id       = 3
       dvt mode     = 0
       dvt vsan     = 3
       tgt pwwn     = 22:00:00:20:37:88:20:ef
       tgt lun      = 0x0
       tgt vsan     = 1
```

Example 49-4 Displays SANTap Session Information

```
switch# show santap module 2 session
Session Information :
       session id   = 1
       host pwwn     = 21:00:00:e0:8b:12:8b:7a
       dvt pwwn     = 50:06:0e:80:03:81:32:36
       dvt lun      = 0x0
       tgt pwwn     = 50:06:0e:80:03:81:32:36
       tgt lun      = 0x0
       adt pwwn     = 33:33:33:33:33:33:33:00
       adt lun      = 0x0
       aci pwwn     = 22:22:22:22:22:22:22:22
       cvt pwwn     = 23:4f:00:0d:ec:09:3c:02
       num ranges   = 0
       session state = 5
       redirect mode = 0
       mrl requested 1
       MRL : vsan 8 RegionSize 4806720, DiskPWWN 0x234f000dec093c02, DiskLun 0x 1,
startLBA 1
```

```
        pwl requested 1
        PWL : type 2, UpdatePol 2, RetirePolicy 4, pwl_start 1

        iol requested 0
```

Example 49-5 Displays SANTap AVT Information

```
switch# show santap module 2 avt

 AVT Information :
        avt pwwn     = 2a:4b:00:05:30:00:22:25
        avt nwwn     = 2a:60:00:05:30:00:22:25
        avt id       = 12
        avt vsan     = 4
        avt if_index = 0x1080000
        hi pwwn      = 21:00:00:e0:8b:07:61:aa
        tgt pwwn     = 22:00:00:20:37:88:20:ef
        tgt vsan     = 1
```

Example 49-6 Displays SANTap AVT LUN Information

```
switch# show santap module 2 avtlun

 AVT LUN Information :
        avt pwwn     = 2a:4b:00:05:30:00:22:25
        avt lun      = 0x0
        xmap id      = 16
        avt id       = 12
        tgt lun      = 0x0
```

Example 49-7 Displays SANTap Remote Virtual Terminal Information

```
switch# show santap module 2 rvt

 RVT Information :
        rvt pwwn     = 2a:61:00:05:30:00:22:25
        rvt nwwn     = 2a:62:00:05:30:00:22:25
        rvt id       = 17
        rvt vsan     = 4
        rvt if_index = 0x1080000
```

Example 49-8 Displays SANTap Remote Virtual Terminal LUN Information

```
switch# show santap module 2 rvtlun

 RVT LUN Information :
        rvt pwwn     = 2a:61:00:05:30:00:22:25
        rvt lun      = 0x0
        xmap id      = 22
        rvt id       = 17
        app pwwn     = 22:00:00:20:37:39:b1:00
        app lun      = 0x0
        app vsan     = 1
```

Removing Appliance-Generated Entities

An appliance might terminate its SANTap application without removing generated entities on the MDS switch. This section describes how to remove these entities using the CLI on the MDS switch.

This section includes the following topics:

Removing AVTs and AVT LUNs

Occasionally the AVT and AVT LUN configuration remains after a SANTap application terminates. To remove AVTs and AVT LUNs, follow these steps:

	Command	Purpose
Step 1	`switch# show santap module 2 avt`	Displays the AVT pWWNs.
	`switch# show santap module 2 avtlun`	Displays the AVT pWWNs and LUNs
Step 2	`switch# clear santap module 2 avt 2a:4b:00:05:30:00:22:25 lun 0x0`	Removes a LUN from the AVT.
	`switch# clear santap module 2 avt 2a:4b:00:05:30:00:22:25`	Removes the AVT. **Note** You can remove the AVT only after all the LUNs are removed.

Removing SANTap Sessions

Occasionally a SANTap session continues after a SANTap application terminates. To remove a SANTap session, follow these steps:

	Command	Purpose
Step 1	`switch# show santap module 2 session`	Displays SANTap session information on the SSM in slot 2.
Step 2	`switch# clear santap module 2 session 1`	Removes SANTap session 1 on the SSM in slot 2.

Removing Initiator-Target-LUNs

The initiator-target-LUN (ITL) triplet identifies a LUN loaded on a DVT. Occasionally the ITL configuration remains after a SANTap application terminates. To remove all LUNs for an ITL triplet, follow these steps:

	Command	Purpose
Step 1	`switch# show santap module 2 dvtlun`	Displays the target and host pWWNs for the ITLs on the SSM in slot 2.
Step 2	`switch# clear santap module 2 itl target-pwwn 22:00:00:20:37:88:20:ef host-pwwn 22:00:00:20:37:88:20:ef`	Removes an IT L on the SSM in slot 2.

Default Settings

Table 49-1 lists the default settings for SANTap parameters.

Table 49-1 Default SANTap Parameters

Parameters	Default
SANTap feature	Disabled.
DVT IO timeout	10 seconds.
DVT LUN size handling flag	0 (disabled).

CHAPTER 50

Configuring NASB

The Storage Services Module (SSM) supports Network-Accelerated Serverless Backup (NASB).

For licensing details, see Chapter 3, "Obtaining and Installing Licenses."

This chapter includes the following sections:

About NASB

Data movement in the fabric uses considerable processor cycles, which can cause client applications to slow down noticeably. Offloading data movement operations to a media server allows the client applications to run normally even during a backup operation. Media servers can further offload the data movement operation to NASB devices, which allows the media server to focus on the coordination functions needed to complete the backup.

Most backups performed today are server-free. In server-free backups, the application server is not involved in moving the data. The data can be moved by either a media server or a NASB device.

When the media server is the data mover, it moves the data between the disks and the tapes. The backup application runs on both the client device and the media server. However, the backup application in the client device performs minimal tasks for the backup operation.

The media server performs the following backup operations:

- Manages disks as well as one or more tape backup devices.
- Contacts the client devices to retrieve the list of logical blocks that need to be backed up.
- Performs data movement from disk to tape media based on the logical block list provided by the client device.

The backup application in the client device maps the data to be backed up and creates the logical block list associated with the data. The movement of data from the physical disks to the backup device (tape) is not performed by the client device. This reduces substantial load on the client device.

Note The media server, disk, and tape can be located anywhere in the fabric.

An example configuration is shown in Figure 50-1. The media server moves the data directly between the storage disks and the tape devices during backups.

Figure 50-1 Example Configuration with Media Server as Data Mover

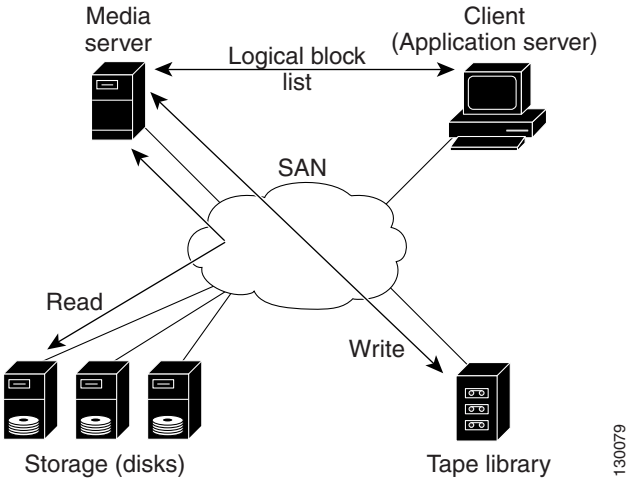

When the NASB is the data mover, it moves the data between the disks and the tapes. The NASB device is a SCSI target device capable of handling SCSI Extended Copy (XCOPY) commands as well as a SCSI initiator device capable of issuing READ/WRITE commands to disks and other backup media, such as tapes. See Figure 50-2.

Figure 50-2 Example Configuration with NASB Device as Data Mover

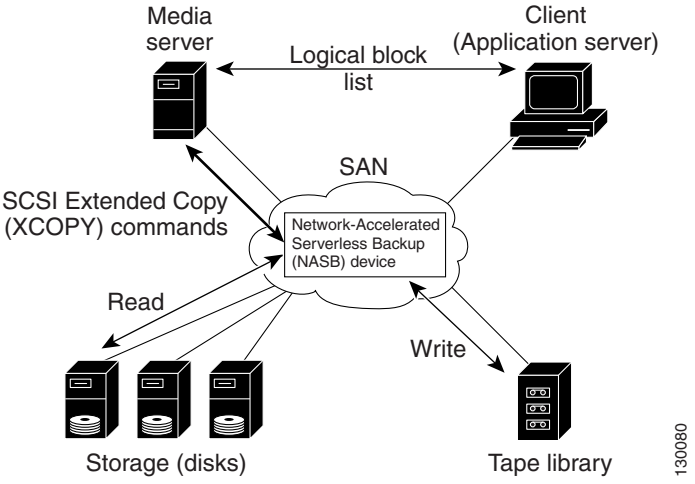

The task of managing and preparing the source and destination targets is performed by the media server. For example, if the destination is a tape library, the media server issues commands to load and unload the correct tape and position of the tape write head at the correct offset within the tape.

Configuring NASB

Network-Accelerated Serverless Backup (NASB) can be enabled on an entire SSM or it can be enabled on one or more groups of four ports on an SSM. Enabling NASB on interfaces has the following restrictions:

- The fewest number of interfaces that you can enable is four. You can specify fc1 through fc4 but not fc1 through fc2.
- The first interface in the group must be 1, 5, 9, 13, 17, 21, 25, or 29. You can specify fc5 through fc8 but not fc7 through fc10.
- The groups of four interfaces do not need to be consecutive. You can specify fc1 through fc8 and fc17 through fc20.

To configure the NASB feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# ssm enable feature nasb module 4`	Enables the NASB application on the entire SSM in slot 4.
	`switch(config)# no ssm enable feature nasb module 4`	Disables the NASB application on the entire SSM in slot 4.
	`switch(config)# no ssm enable feature nasb force module 4`	Forces the switch to disable the NASB application on the entire SSM in slot 4.
Step 3	`switch(config)# ssm enable feature nasb interface fc 4/1 - 4`	Enables the NASB application on ports 1 through 4 on the SSM in slot 4. **Note** Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29.
	`switch(config)# no ssm enable feature nasb interface fc 4/1 - 4`	Disables the NASB application on ports 1 through 4 on the SSM in slot 4.
	`switch(config)# no ssm enable feature nasb force interface fc 4/1 - 4`	Forces the switch to disable the NASB application on ports 1 through 4 on the SSM in slot 4.

	Command	Purpose
Step 4	`switch(config)# nasb module 4 vsan 10`	Enables NASB on the SSM in slot 4 and on VSAN 10 for a single target LUN. By default, the LUN is a Direct Access Device (Peripheral Device Type = 0x00).
	`switch(config)# nasb module 4 vsan 10 control`	Enables NASB on the SSM in slot 4 and on VSAN 10 for a single target LUN that is a Storage Array Controller (Peripheral Device Type = 0x0C).
	`switch(config)# nasb module 4 vsan 10 multiple`	Enables NASB on the SSM in slot 4 and on VSAN 10 for up to 10 target LUNs that are Direct Access Devices (Peripheral Device Type = 0x00). **Note** Use the **multiple** option for multi-streaming (multiple backup sessions) on a single virtual target for Veritas NetBackup.
	`switch(config)# nasb module 4 vsan 10 control multiple`	Enables NASB on the SSM in slot 4 and on VSAN 10 for Storage Array Controller (Peripheral Device Type = 0x0C) and up to 10 target LUNs.
	`switch(config)# nasb module 4 vsan 10 multiple control`	Enables NASB on the SSM in slot 4 and on VSAN 10 for up to 10 target LUNs and Storage Array Controller (Peripheral Device Type = 0x0C).
Step 5	`switch(config)# no nasb module 4 vsan 10`	Disables NASB.

Note You cannot simultaneously configure the intelligent services SANTap and NASB on a single SSM.

NASB Target Rediscovery

You can initiate a rediscovery of a target device (disk or tape) if the configuration on the target side has changed without generating an RSCN in the fabric, such as a change in the access list or LUN-mapping on the target. Use the following step to initiate target device rediscovery:

	Command	Purpose
Step 1	`switch# nasb rediscover module 2 vsan 9 target-pwwn 20:02:00:a0:b8:16:a1:5f` `nasb rediscovery initiated`	Initiates a rediscovery of a target device for the SSM in slot 2.

Displaying NASB Information

Use the **show nasb** command to display information about NASB (see Example 50-1 to Example 50-4).

Example 50-1 Displays NASB Information

```
switch# show nasb
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 50-2 Displays NASB Information for a Specific Module

```
switch# show nasb module 3
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 50-3 Displays NASB Information for a Specific Module for a VSAN

```
switch# show nasb module 3 vsan 2
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 50-4 Displays NASB Information for a Specific VSAN

```
switch# show nasb vsan 1
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
```

```
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
```

Default Settings

Table 50-1 lists the default settings for NASB parameters.

Table 50-1 Default NASB Parameters

Parameters	Default
NASB feature	Disabled.

PART 8

Network and Switch Monitoring

CHAPTER 51

Configuring RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or later software.

This chapter includes the following sections:

About RMON

All switches in the Cisco MDS 9000 Family support the following RMON functions (defined in RFC 2819):

- Alarm—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Event—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for agent and management information.

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for information on an SNMP-compatible network management station.

See the "About SNMP Security" section on page 32-1 for SNMP security-related CLI configurations.

Configuring RMON

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or an SNMP-compatible network management station.

Tip We recommend an additional, generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Note You must also configure SNMP on the switch to access RMON MIB objects.

RMON Alarm Configuration

You can set an alarm on any MIB object. The specified MIB must be an existing SNMP MIB object in standard dot notation (1.3.6.1.2.1.2.2.1.14.16777216 for ifInOctets.16777216).

Use one of the following options to specify the interval to monitor the MIB variable (ranges from 1 to 4294967295 seconds):

- Use the **delta** option to test the change between samples of a MIB variable.
- Use the **absolute** option to test each MIB variable directly.
- Use the **delta** option to test any MIB objects that are counters.

The range for the **rising threshold** and **falling threshold** values is -2147483647 to 2147483647.

Caution The **falling threshold** must be less than the **rising threshold**.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.
- The owner of the alarm.

To enable RMON alarms, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# rmon alarm 20 1.3.6.1.2.1.2.2.1.14.16777216 2900 delta rising-threshold 15 1 falling-threshold 0 owner test`	Configures RMON alarm number 20 to monitor the 1.3.6.1.2.1.2.2.1.14.16777216 once every 900 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the value shows a MIB counter increase of 15 or more, the software triggers an alarm. The alarm in turn triggers event number 1, which is configured with the RMON event command. Possible events can include a log entry or an SNMP trap. If the MIB value changes by 0, the alarm is reset and can be triggered again.
	`switch(config)# no rmon alarm 2`	Deletes the specified entry from the alarm table.

RMON Event Configuration

To enable RMON events, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# rmon event 2 log trap eventtrap description CriticalErrors owner Test2`	Creates RMON event number 2 to define CriticalErrors and generates a log entry when the event is triggered by the alarm. The user Test2 owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.
	`switch(config)# no rmon event 5`	Deletes an entry from the RMON event table.

RMON Verification

Use the **show rmon** and **show snmp** commands to display configured RMON and SNMP information (see Example 51-1 and 51-3).

Example 51-1 Displays Configured RMON Alarms

```
switch# show rmon alarms
Alarm 1 is active, owned by admin
 Monitors 1.3.6.1.2.1.2.2.1.16.16777216 every 1 second(s)
 Taking delta samples, last value was 0
 Rising threshold is 1, assigned to event 0
 Falling threshold is 0, assigned to event 0
 On startup enable rising or falling alarm
```

Example 51-2 Displays Configured RMON High Capacity Alarms

```
switch# show rmon hcalarms
High Capacity Alarm 10 is active, owned by Testuser
 Monitors 1.3.6.1.2.1.31.1.1.1.6.16785408 every 300 second(s)
 Taking absolute samples, last value was 0 (valuePositive)
 Rising threshold low is 4294967295 & high is 15 (valuePositive)
 Rising threshold assigned to event 1
 Falling threshold low is 0 & high is 0 (valueNotAvailable)
 Falling threshold assigned to event 0
 On startup enable rising alarm
 Number of Failed Attempts is 0
```

Note High capacity RMON alarms can be configured using the CISCO-HC-ALARM-MIB. See the *Cisco MDS 9000 Family MIB Quick Reference*.

Example 51-3 Displays Configured RMON Events

```
switch# show rmon events
Event 2 is active, owned by Test2
 Description is CriticalErrors
 Event firing causes log and trap to community eventtrap, last fired 0
```

```
Event 500 is active, owned by admin
 Description is
 Event firing causes log, last fired 138807208
```

Default Settings

Table 51-1 lists the default settings for all RMON features in any switch.

Table 51-1 Default RMON Settings

Parameters	Default
RMON alarms	Disabled.
RMON events	Disabled.

CHAPTER 52

Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic (see the"Cisco Fabric Analyzer" section on page 58-4.

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is non-intrusive and does not affect switching of network traffic for any SPAN source ports (see Figure 52-1).

Figure 52-1 SPAN Transmission

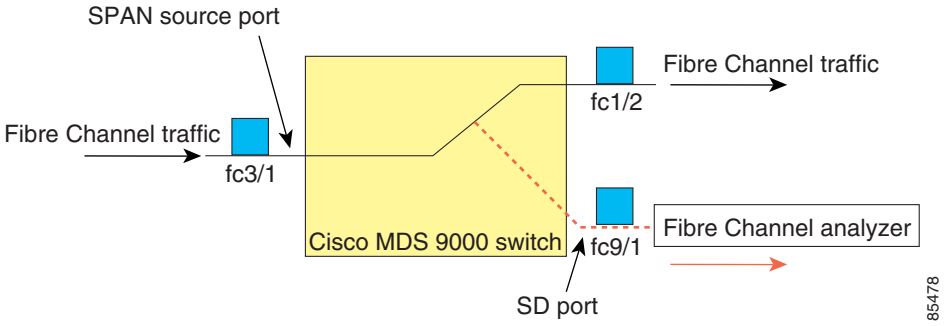

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see Figure 52-2).

Figure 52-2 SPAN Traffic from the Ingress Direction

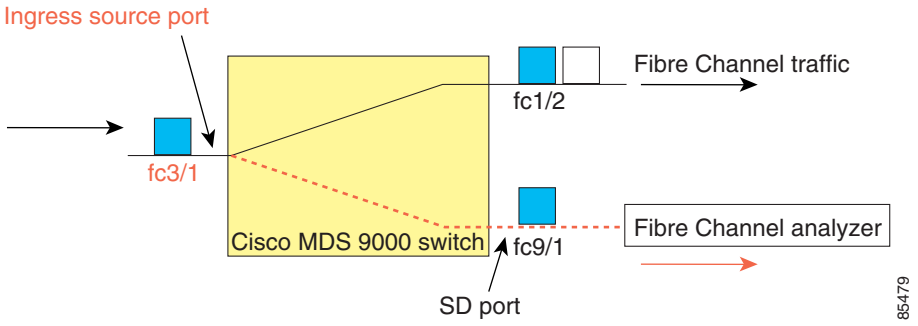

- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 52-3).

Figure 52-3 SPAN Traffic from Egress Direction

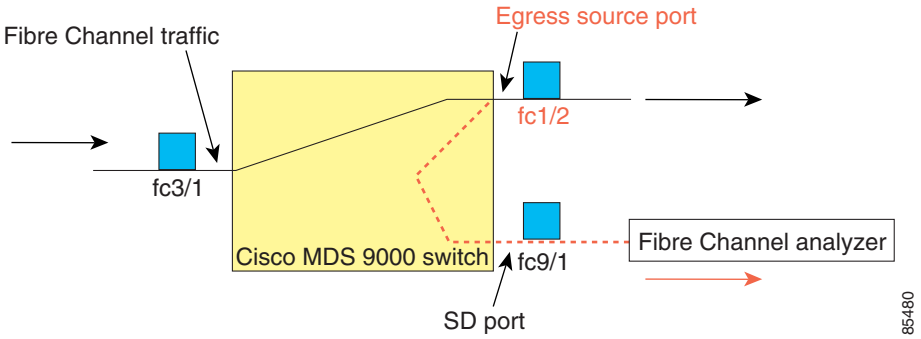

IPS Source Ports

SPAN capabilities are available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.

Note You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
  - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
  - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
  - All ports in the PortChannel are included and spanned as sources.
  - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces:
  - iSCSI interfaces
  - FCIP interfaces

VSAN as a Source

When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.
- Interfaces are only included as sources when the port VSAN matches the source VSAN. Figure 52-4 displays a configuration using VSAN 2 as a source:
  - All ports in the switch are in VSAN 1 except fc1/1.
  - Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
  - VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 52-4 VSAN as a Source

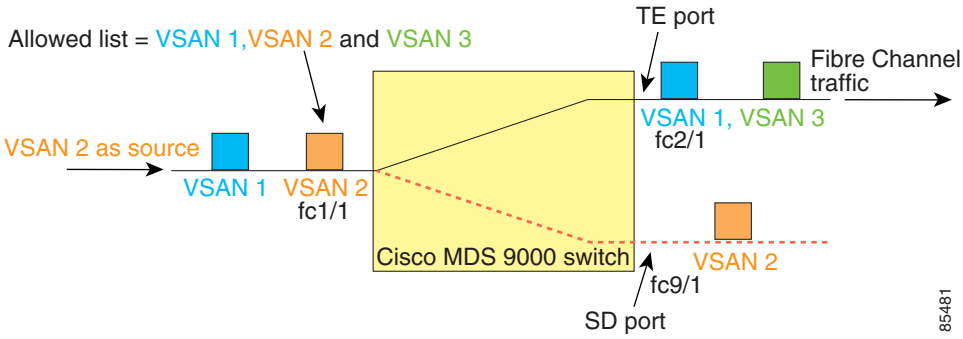

For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 because the port VSAN does not match VSAN 1.

  See the "Configuring an Allowed-Active List of VSANs" section on page 15-6 or the "About Port VSAN Membership" section on page 19-7.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.

Tip A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session (see Figure 52-4). Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.
- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.

Note If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode using the **switchport mode** command.

Guidelines to Configure SPAN

The following guidelines apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit (see the "32-Port Switching Module Configuration Guidelines" section on page 12-2).
- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.

Configuring SPAN

To monitor network traffic using SD ports, follow these steps:

Step 1 Configure the SD port.

Step 2 Attach the SD port to a specific SPAN session.

Step 3 Monitor network traffic by adding source interfaces to the session.

To configure an SD port for SPAN monitoring, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc9/1`	Configures the specified interface.
Step 3	`switch(config-if)# switchport mode SD`	Configures the SD port mode for interface fc9/1.
Step 4	`switch(config-if)# switchport speed 1000`	Configures the SD port speed to 1000 Mbps.
Step 5	`switch(config-if)# no shutdown`	Enables traffic flow through this interface.

To configure a SPAN session, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Configures the specified SPAN session (1). If the session does not exist, it is created.
	`switch(config)# no span session 1`	Deletes the specified SPAN session (1).
Step 3	`switch(config-span)# destination interface fc9/1`	Configures the specified destination interface (fc 9/1) in a session.
	`switch(config-span)# no destination interface fc9/1`	Removes the specified destination interface (fc 9/1).
Step 4	`switch(config-span)# source interface fc7/1`	Configures the source (fc7/1) interface in both directions. **Note** The Cisco MDS 9124 Fabric Switch does not support bi-directional SPAN sessions (Rx and Tx)
	`switch(config-span)# no source interface fc7/1`	Removes the specified destination interface (fc 7/1) from this session.

	Command	Purpose
Step 5	`switch(config-span)# source interface sup-fc0`	Configures the source interface (sup-fc0) in the session.
	`switch(config-span)# source interface fc1/5 - 6, fc2/1 -3`	Configures the specified interface ranges in the session.
	`switch(config-span)# source vsan 1-2`	Configures source VSANs 1 and 2 in the session.
	`switch(config-span)# source interface port-channel 1`	Configures the source PortChannel (port-channel 1).
	`switch(config-span)# source interface fcip 51`	Configures the source FCIP interface in the session.
	`switch(config-span)# source interface iscsi 4/1`	Configures the source iSCSI interface in the session.
	`switch(config-span)# source interface svc1/1 tx traffic-type initiator`	Configures the source SVC interface in the Tx direction for an initiator traffic type.
	`switch(config-span)# no source interface port-channel 1`	Deletes the specified source interface (port-channel 1).

To configure a SPAN filter, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Configures the specified session (1).
Step 3	`switch(config-span)# source interface fc9/1 tx`	Configures the source fc9/1 interface in the egress (Tx) direction.
	`switch(config-span)# source filter vsan 1-2`	Configures VSANs 1 and 2 as session filters.
	`switch(config-span)# source interface fc7/1 rx`	Configures the source fc7/1 interface in the ingress (Rx) direction.

Configuring SPAN for Generation 2 Fabric Switches

Cisco Generation 2 Fabric Switches (such as MDS 9124) support single direction (SPAN session 1) SPAN sessions only.

The following example shows how to configure SPAN sessions for the ingress direction for this switch.

Example 52-1 Configuring a Generation 2 Fabric Switch for Ingress SPAN sessions

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Configures the specified session (1).

	Command	Purpose
Step 3	`switch(config-span)# destination interface fc1/1`	Configures interface fc1/1 as the destination.
Step 4	`switch(config-span)# source interface fc1/2 rx`	Configures the source interface fc1/2 in the ingress direction.

Example 52-2 Configuring a Generation 2 Fabric Switch for Egress SPAN Session

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Configures the specified session (1).
Step 3	`switch(config-span)# destination interface fc1/1`	Configures interface fc1/1 as the destination.
Step 4	`switch(config-span)# source interface fc1/2 tx`	Configures the source interface fc1/2 in the egress direction.

You cannot mix ingress and egress interfaces in the same SPAN session. The SPAN will reject any configuration that mixes Rx ad Tx directions. However, you can specify multiple SPAN source interfaces in a single direction.

Example 52-3 Configuring Cisco MDS 9124 for Multiple SPAN Interfaces in a Single Direction

```
switch(config-span)# span session 1
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 rx
switch(config-span)# source interface fc1/3 rx
```

Generation 2 Fabric Switches support VSAN filters for one VSAN only in the egress direction; this restriction does not apply to the ingress direction. For example, if you have an interface that is a TE port, with an active VSAN of 1 to 5, and you specify a VSAN filter for VSAN 2, then only the traffic on VSAN 2 will be filtered.

```
switch(config-span)# span session 1
switch(config-span)# source filter vsan 2
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 tx
```

However, if you specify the VSAN filter for VSANs 1 to 2, then traffic from all VSANs (1 to 5) is filtered–essentially rendering the filter useless.

```
switch(config-span)# span session 1
switch(config-span)# source filter vsan 1-2
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 tx
```

Suspending and Reactivating SPAN Sessions

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

To temporarily suspend or reactivate a SPAN session filter, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Configures the specified session (1).
Step 3	`switch(config-span)# suspend`	Temporarily suspends the session.
	`switch(config-span)# no suspend`	Reactivates the session.

Note When using Generation 2 Fabric Switches, you cannot create an additional active SPAN session when you already have one.

Encapsulating Frames

The frame encapsulation feature is disabled by default. If you enable the encapsulation feature, all outgoing frames are encapsulated.

The **switchport encap eisl** command only applies to SD port interfaces. If encapsulation is enabled, you see a new line (`Encapsulation is eisl`) in the **show interface** *SD_port_interface* command output.

To encapsulate outgoing frames (optional), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc9/32`	Configures the specified interface.
Step 3	`switch(config-if)# switchport mode SD`	Configures the SD port mode for interface fc9/32.
Step 4	`switch(config-if)# switchport encap eisl`	Enables the encapsulation option for this SD port.
	`switch(config-if)# no switchport encap eisl`	Disables (default) the encapsulation option.

SPAN Conversion Behavior

SPAN features (configured in any prior release) are converted as follows:

- If source interfaces and source VSANs are configured in a given session, then all the source VSANs are removed from that session.

  For example, before Cisco MDS SAN-OS Release 1.0(4):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 10-11
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

Session 1 had both source interfaces and source VSANs before the upgrade. After the upgrade, the source VSANs were removed (rule 1).

- If interface level VSAN filters are configured in source interfaces, then the source interfaces are also removed from the session. If this interface is configured in both directions, it is removed from both directions.

  For example, before Cisco MDS SAN-OS Release 1.0(4):

```
Session 2 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 12
     fc1/6 (vsan 1-20),
   Egress (tx) sources are
     fc1/6 (vsan 1-20),
```

  Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 2 (inactive as no active sources)
   Destination is fc1/9
   No session filters configured
   No ingress (rx) sources
   No egress (tx) sources
```

Note The deprecated configurations are removed from persistent memory once a switchover or a new startup configuration is implemented.

  Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Cisco MDS SAN-OS Release 1.0(4). When upgraded to Cisco MDS SAN-OS Release 1.1(1) the following changes are made:

  - The source VSAN (VSAN 12) is removed (rule 1).
  - The source interface fc1/6 had VSAN filters specified—it is also removed (rule 2).

Monitoring Traffic Using Fibre Channel Analyzers

You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is specially useful in troubleshooting scenarios where traffic disruption changes the problem environment and makes it difficult to reproduce the problem.

Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 as shown in Figure 52-5.

Figure 52-5 Fibre Channel Analyzer Usage Without SPAN

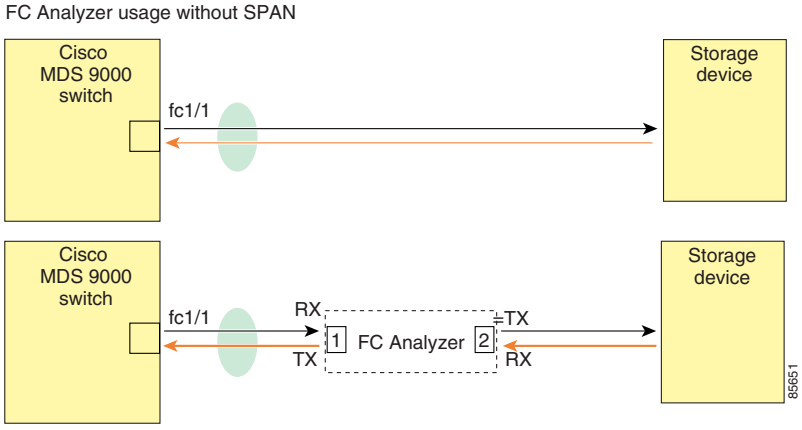

This type of connection has the following limitations:

- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

With SPAN

Using SPAN you can capture the same traffic scenario shown in Figure 52-5 without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer as shown in Figure 52-6.

Figure 52-6 Fibre Channel Analyzer Using SPAN

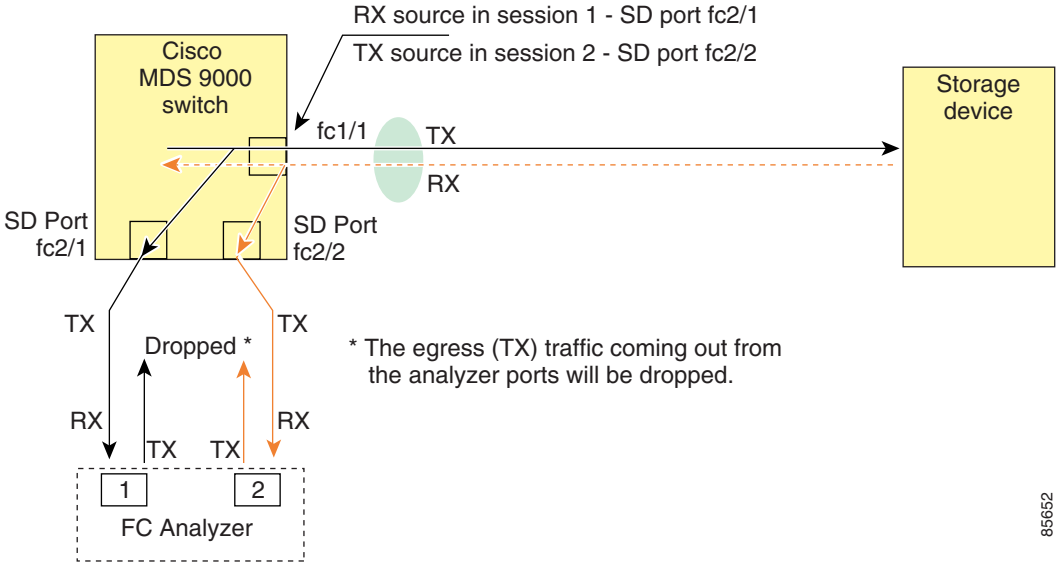

Configuring Fibre Channel Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in Figure 52-6, follow these steps:

Step 1 Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.

Step 2 Configure SPAN on interface fc1/1in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.

Step 3 Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.

Step 4 Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.

To configure SPAN on the source and destination interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Creates the SPAN session 1.
Step 3	`switch(config-span)## destination interface fc2/1`	Configures the destination interface fc2/1.
Step 4	`switch(config-span)# source interface fc1/1 rx`	Configures the source interface fc1/1 in the ingress direction.
Step 5	`switch(config)# span session 2` `switch(config-span)#`	Creates the SPAN session 2.
Step 6	`switch(config-span)## destination interface fc2/2`	Configures the destination interface fc2/2.
Step 7	`switch(config-span)# source interface fc1/1 tx`	Configures the source interface fc1/1 in the egress direction.

Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface as shown in Figure 52-6. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

Figure 52-7 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions. This setup is more advantageous and cost effective than the setup shown in Figure 52-6—it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 52-7 Fibre Channel Analyzer Using a Single SD Port

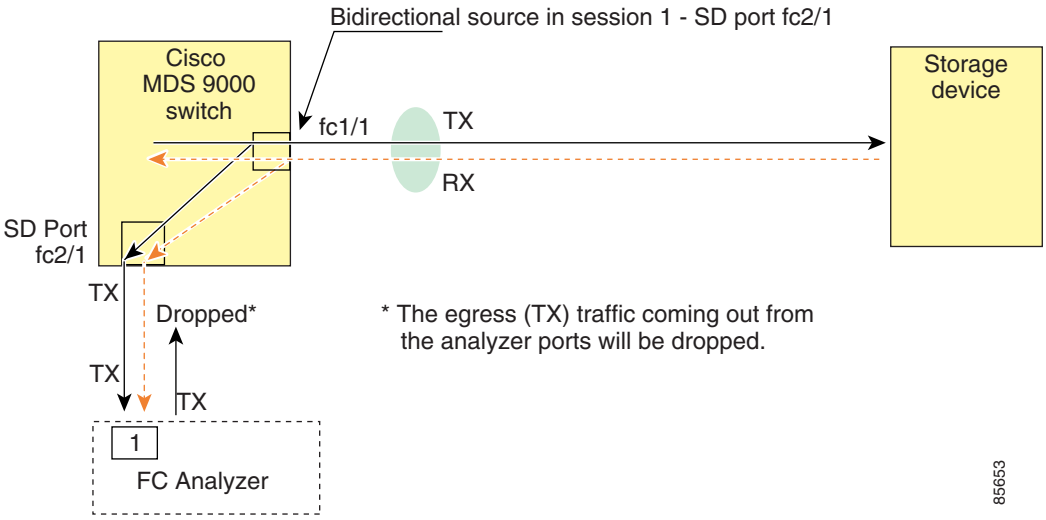

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

To configure SPAN on a single SD port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# span session 1` `switch(config-span)#`	Creates the SPAN session 1.
Step 3	`switch(config-span)## destination interface fc2/1`	Configures the destination interface fc2/1.
Step 4	`switch(config-span)# source interface fc1/1`	Configures the source interface fc1/1 on the same SD port.

Displaying SPAN Information

Use the **show span** command to display configured SPAN information. See Examples 52-4 to 52-7.

Example 52-4 Displays SPAN Sessions in a Brief Format

```
switch# show span session brief
-------------------------------------------------------
 Session  Admin          Oper           Destination
          State          State          Interface
-------------------------------------------------------
 7        no suspend     active         fc2/7
 1        suspend        inactive       not configured
 2        no suspend     inactive       fc3/1
```

Example 52-5 Displays a Specific SPAN Session in Detail

```
switch# show span session 7
Session 7 (active)
   Destination is fc2/7
   No session filters configured
   No ingress (rx) sources
   Egress (tx) sources are
     port-channel 7,
```

Example 52-6 Displays ALL SPAN Sessions

```
switch# show span session
Session 1 (inactive as no destination)
Destination is not specified
   Session filter vsans are 1
   No ingress (rx) sources
   No egress (tx) sources
Session 2 (active)
   Destination is fc9/5
   No session filters configured
   Ingress (rx) sources are
      vsans 1
   No egress (tx) sources
Session 3 (admin suspended)
   Destination is not configured
   Session filter vsans are 1-20
   Ingress (rx) sources are
     fc3/2, fc3/3, fc3/4, fcip 51,
     port-channel 2, sup-fc0,
   Egress (tx) sources are
     fc3/2, fc3/3, fc3/4, sup-fc0,
```

Example 52-7 Displays an SD Port Interface with Encapsulation Enabled

```
switch# show int fc9/32
fc9/32 is up
    Hardware is Fibre Channel
    Port WWN is 22:20:00:05:30:00:49:5e
    Admin port mode is SD
    Port mode is SD
    Port vsan is 1
    Speed is 1 Gbps
    Receive Buffer Size is 2112
    Encapsulation is eisl <---------- Displays the enabled encapsulation status
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes, 0 discards
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes, 0 discards
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

Remote SPAN

Note Remote SPAN is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a Cisco MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for those SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see Figure 52-8):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 52-8 RSPAN Transmission

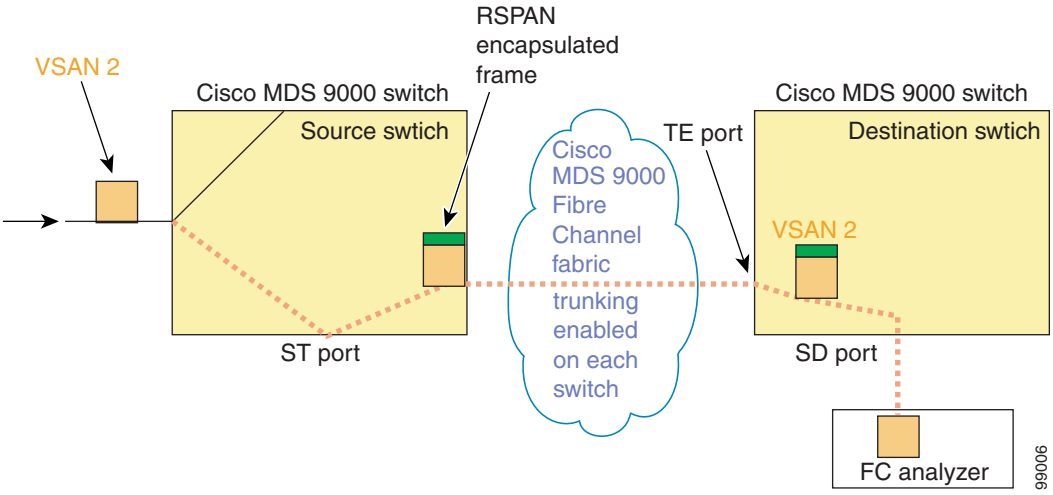

Advantages to Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Is compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

An FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to an ST port in the source switch and map the same FC tunnel to an SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as an RSPAN tunnel (see Figure 52-9).

Figure 52-9 FC and RSPAN Tunnel

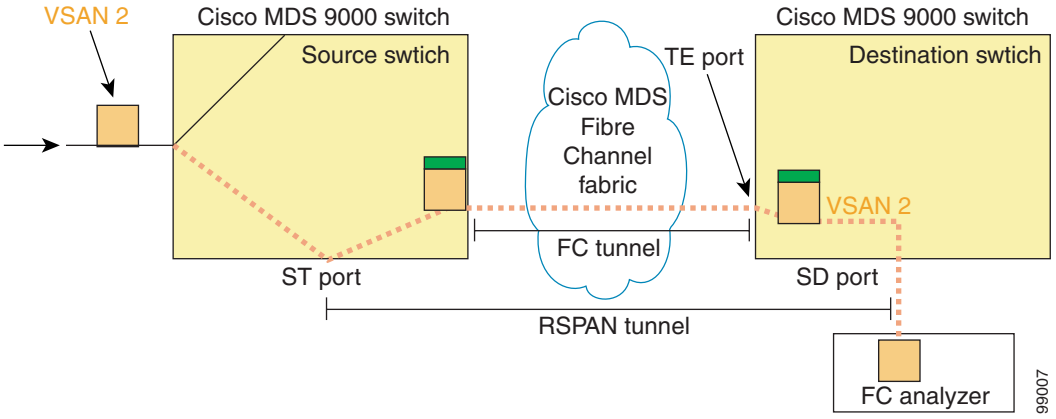

Guidelines to Configure RSPAN

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it is dropped.
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented:
  - Trunking must be enabled (enabled by default).
  - VSAN interface must be configured.
  - The Fibre Channel tunnel feature must be enabled (disabled by default).
  - IP routing must be enabled (disabled by default).

Note If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface

  See Chapter 43, "Configuring IP Services."

ST Port Characteristics

ST ports have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.
- One ST port can only be bound to one FC tunnel.

- ST ports cannot be used for any purpose other than to carry RSPAN traffic.
- ST ports cannot be configured using Storage Services Modules (SSMs).

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.

Note Besides the source and destination switches, the VSAN must also be configured in each Cisco MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

Step 1 Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation.

Step 2 Enable the FC tunnel in each switch in the end-to-end path of the tunnel.

Step 3 Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface's IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port.

Step 4 Configure SD ports for SPAN monitoring in the destination switch (Switch D).

Step 5 Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel.

Step 6 Create an RSPAN session in the source switch (in Switch S) to monitor network traffic.

RSPAN Configuration Example

This section provides a RSPAN configuration example using the procedure defined in the previous section.

Configuration in the Source Switch

This section identifies the tasks that must be performed in the source switch (Switch S):

- Configuring the SD Port, page 52-24
- Mapping the FC Tunnel, page 52-25

Creating VSAN Interfaces

Figure 52-10 depicts a basic FC tunnel configuration.

Figure 52-10 FC Tunnel Configuration

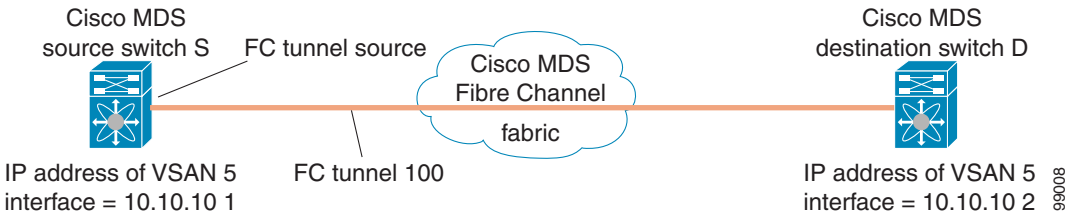

Note This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the source switch for the scenario in Figure 52-10, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# interface vsan 5` `switchS(config-if)#`	Configures the specified VSAN interface (VSAN 5) in the source switch (switch S).
Step 3	`switchS(config-if)# ip address 10.10.10.1 255.255.255.0`	Configures the IPv4 address and subnet for the VSAN interface 5 in the source switch (switch S).
Step 4	`switchS(config-if)# no shutdown`	Enables traffic flow through this interface.

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# fc-tunnel enable`	Enables the FC tunnel feature (disabled by default).

Note Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Initiating the FC Tunnel

To initiate the FC tunnel in the source switch for the scenario in Figure 52-10, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# interface fc-tunnel 100` `switchS(config-if)#`	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.
Step 3	`switchS(config-if)# source 10.10.10.1`	Maps the IPv4 address of the source switch (switch S) to the FC tunnel (100).
Step 4	`switchS(config-if)# destination 10.10.10.2`	Maps the IPv4 address of the destination switch (switch D) to the FC tunnel (100).
Step 5	`switchS(config-if)# no shutdown`	Enables traffic flow through this interface.

Tip The interface cannot be operationally up until the FC tunnel mapping is configured in the destination switch.

Configuring the ST Port

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes an RSPAN tunnel once the binding and mapping is complete.

Figure 52-11 depicts a basic FC tunnel configuration.

Figure 52-11 Binding the FC Tunnel

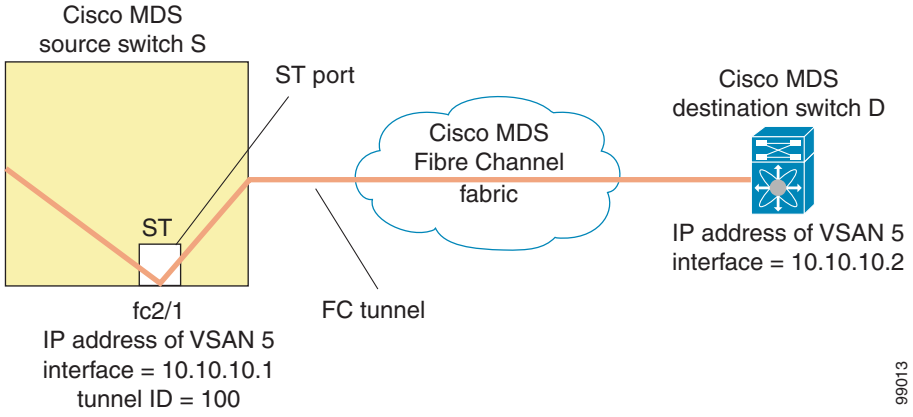

Note ST ports cannot be configured using Storage Services Modules (SSMs).

To configure an ST port for the scenario in Figure 52-11, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# interface fc2/1`	Configures the specified interface.

	Command	Purpose
Step 3	`switchS(config-if)# switchport mode ST`	Configures the ST port mode for interface fc2/1.
Step 4	`switchS(config-if)# switchport speed 2000`	Configures the ST port speed to 2000 Mbps.
Step 5	`switchS(config-if)# rspan-tunnel interface fc-tunnel 100`	Associates and binds the ST port with the RSPAN tunnel (100).
Step 6	`switchS(config-if)# no shutdown`	Enables traffic flow through this interface.

Configuring an RSPAN Session

A RSPAN session is similar to a SPAN session, with the destination interface being an RSPAN tunnel.

To configure an RSPAN session in the source switch for the scenario in Figure 52-11, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# span session 2` `switchS(config-span)#`	Configures the specified SPAN session (2). If the session does not exist, it is created. The session ID ranges from 1 to 16.
Step 3	`switchS(config-span)# destination interface fc-tunnel 100`	Configures the specified RSPAN tunnel (100) in a session.
Step 4	`switchS(config-span)# source interface fc1/1`	Configures the source interface (fc1/1) for this session and spans the traffic from interface fc1/1 to RSPAN tunnel 100.

Configuration in All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

- Configuring VSAN Interfaces, page 52-22
- Enabling FC Tunnels, page 52-23
- Enabling IP Routing, page 52-23

Configuring VSAN Interfaces

Figure 52-12 on page 52-24 depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).

Note This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in Figure 52-12 on page 52-24, follow these steps:

	Command	Purpose
Step 1	`switchD# config t`	Enters configuration mode.
Step 2	`switchD(config)# interface vsan 5` `switchD(config-if)#`	Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D).

	Command	Purpose
Step 3	`switchD(config-if)# ip address 10.10.10.2 255.255.255.0`	Configures the IPv4 address and subnet for the VSAN interface in the destination switch (Switch D).
Step 4	`switchD(config-if)# no shutdown`	Enables traffic flow to administratively allow traffic (provided the operational state is up).

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# fc-tunnel enable`	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.

Note Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric (see the "Enabling IPv4 Routing" section on page 43-7). This step is required to set up the FC tunnel.

Configuration in the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

- Configuring VSAN Interfaces, page 52-23
- Configuring the SD Port, page 52-24
- Mapping the FC Tunnel, page 52-25

Configuring VSAN Interfaces

Figure 52-12 on page 52-24 depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).

Note This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in Figure 52-12, follow these steps:

	Command	Purpose
Step 1	`switchD# config t`	Enters configuration mode.
Step 2	`switchD(config)# interface vsan 5` `switchD(config-if)#`	Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D).

	Command	Purpose
Step 3	`switchD(config-if)# ip address 10.10.10.2 255.255.255.0`	Configures the IPv4 address and subnet for the VSAN interface in the destination switch (Switch D).
Step 4	`switchD(config-if)# no shutdown`	Enables traffic flow to administratively allow traffic (provided the operational state is up).

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# fc-tunnel enable`	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.

Note Be sure to enable this feature in each switch in the end-to-end path in the tunnel.

Configuring the SD Port

The SD port in the destination switch enables the FC analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. Figure 52-12 depicts an RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 52-12 RSPAN Tunnel Configuration

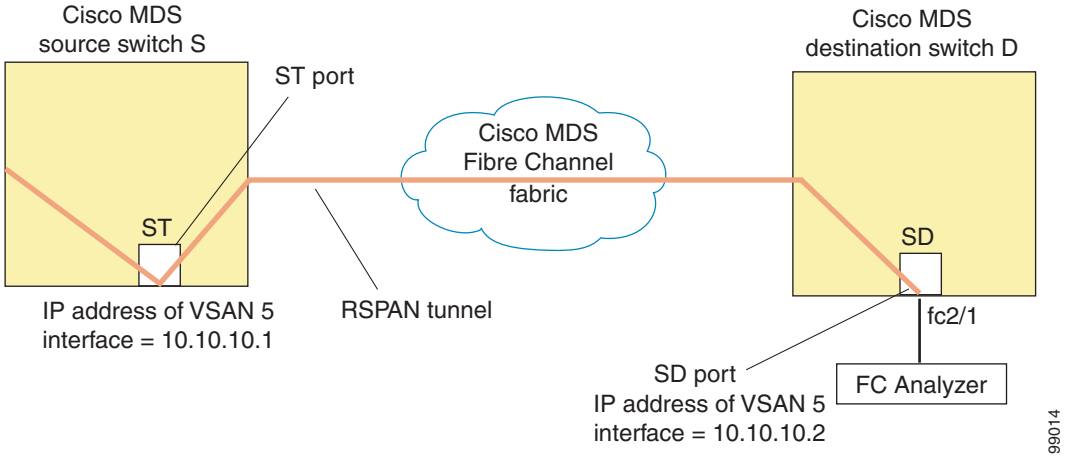

Note SD ports cannot be configured using Storage Services Modules (SSMs).

To configure an SD port for the scenario in Figure 52-12, follow these steps:

	Command	Purpose
Step 1	`switchD# config t`	Enters configuration mode.
Step 2	`switchD(config)# interface fc2/1`	Configures the specified interface.
Step 3	`switchD(config-if)# switchport mode SD`	Configures the SD port mode for interface fc2/1.
Step 4	`switchD(config-if)# switchport speed 2000`	Configures the SD port speed to 2000 Mbps.
Step 5	`switchD(config-if)# no shutdown`	Enables traffic flow through this interface.

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see Figure 52-13).

Figure 52-13 FC Tunnel Configuration

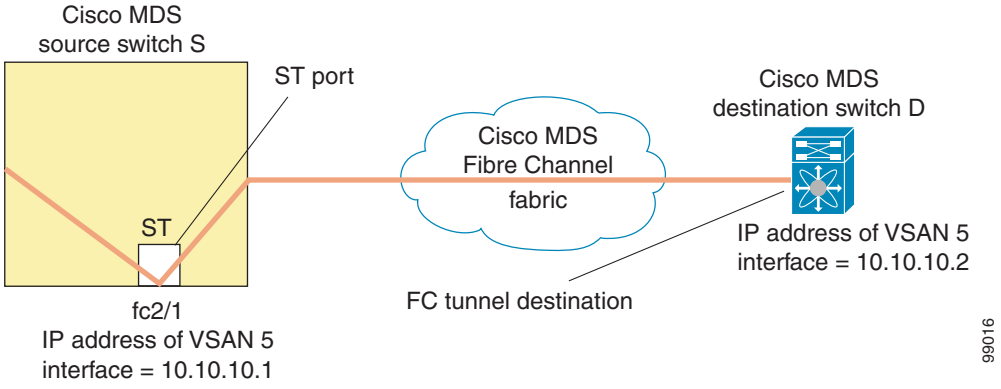

To terminate the FC tunnel in the destination switch for the scenario in Figure 52-13, follow these steps:

	Command	Purpose
Step 1	`switchD# config t`	Enters configuration mode.
Step 2	`switchD(config)# fc-tunnel tunnel-id-map 100 interface fc2/1`	Terminates the FC tunnel (100) in the destination switch (switch D). The tunnel ID range is from 1 to 255.

Explicit Paths

You can specify an explicit path through the Cisco MDS Fibre Channel fabric (source-based routing), using the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the FC tunnel to always take one path to the destination switch. The software then uses this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths are available. In an RSPAN situation, you can specify the explicit path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see Figure 52-14).

Figure 52-14 Explicit Path Configuration

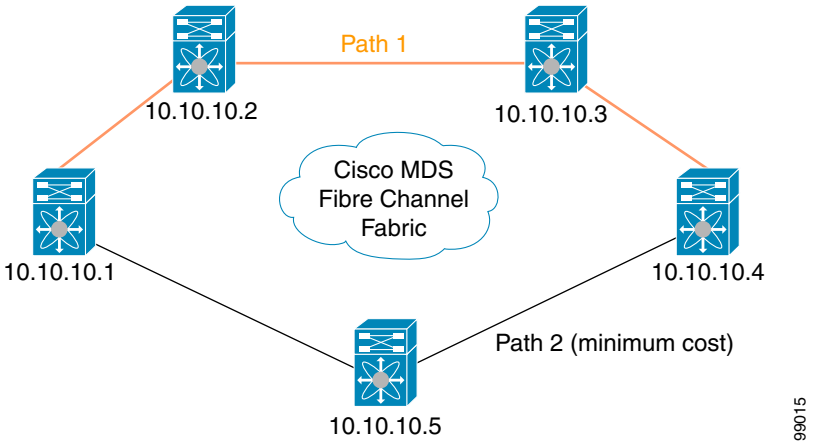

The explicit path must be created in the source switch. To configure an explicit path, you must first create the path and then configure the use of any one path. If an explicit path is not configured, the minimum cost path is used by default. If an explicit path is configured and is functioning, the specified path is used.

To create an explicit path for the scenario in Figure 52-14, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# fc-tunnel explicit-path Path1` `switch(config-explicit-path)#`	Places you at the explicit path prompt for the path named Path 1.

Step 3	`switchS(config-explicit-path)# next-address 10.10.10.2 strict`
	`switchS(config-explicit-path)# next-address 10.10.10.3 strict`
	`switchS(config-explicit-path)# next-address 10.10.10.4 strict`
		Specifies that the next hop VSAN interface IPv4 addresses and the previous hops specified in the explicit path do not require direct connection.

Step 4	`switchS(config)# fc-tunnel explicit-path Path2` `switch(config-explicit-path)#`	Places you at the explicit path prompt for Path2.
Step 5	`switchS(config-explicit-path)# next-address 10.10.10.5 strict` `switchS(config-explicit-path)# next-address 10.10.10.4 strict`	Specifies that the next hop VSAN interface IPv4 addresses and the previous hops specified in the explicit path do not require direct connection.
Step 6	`switchS(config)# fc-tunnel explicit-path Path3` `switch(config-explicit-path)#`	Places you at the explicit path prompt for Path3.
Step 7	`switchS(config-explicit-path)# next-address 10.10.10.3 loose`	Configures a minimum cost path in which the 10.10.10.3 IPv4 address exists. **Note** In Figure 52-14, Path 3 is the same as Path 1—10.10.10.3 exists in Path 1. Using the **loose** option, you can achieve the same effect with one command instead of issuing three commands (using the **strict** option) in Step 3.

To reference the explicit path, follow these steps:

	Command	Purpose
Step 1	`switchS# config t`	Enters configuration mode.
Step 2	`switchS(config)# interface fc-tunnel 100`	References the tunnel ID for Path1.
Step 3	`switchS(config)# explicit-path Path1`	Links Path1 to the tunnel ID.

This configuration explicitly specifies Path 1 to be used for the RSPAN traffic. Refer to RFC 3209 for further details on explicit paths and source based routing.

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. Figure 52-15 shows an RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions.

Figure 52-15 Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic

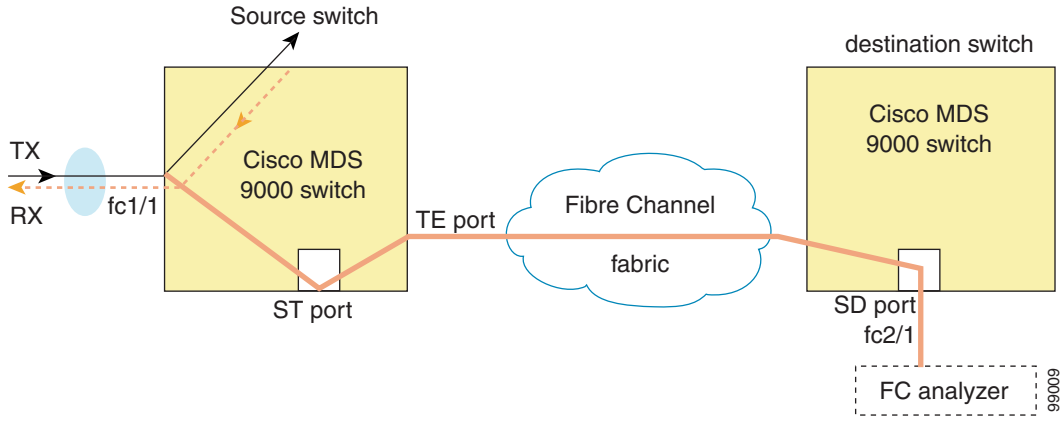

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Sample Scenarios

Note RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. An RSPAN tunnel is configured as a destination interface for the SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel (see Figure 52-16).

Figure 52-16 RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel

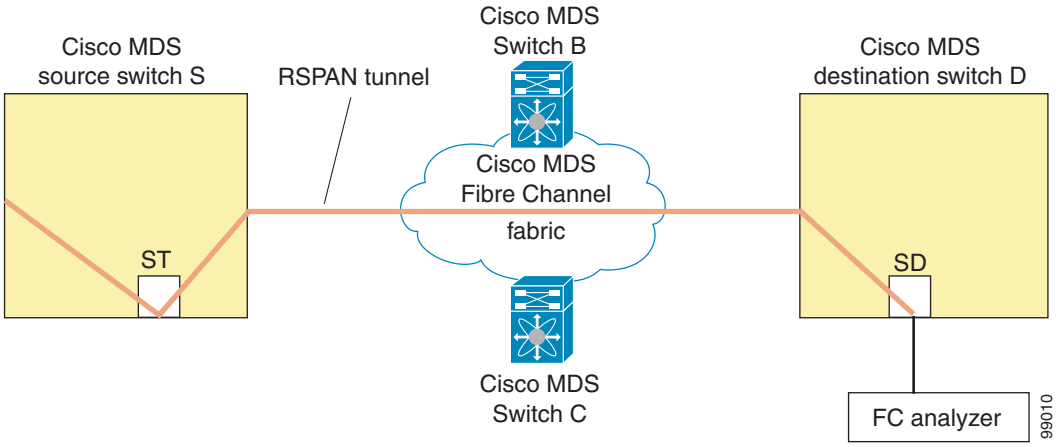

Single Source with Multiple RSPAN Tunnels

Figure 52-17 displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for troubleshooting purposes.

Figure 52-17 RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels

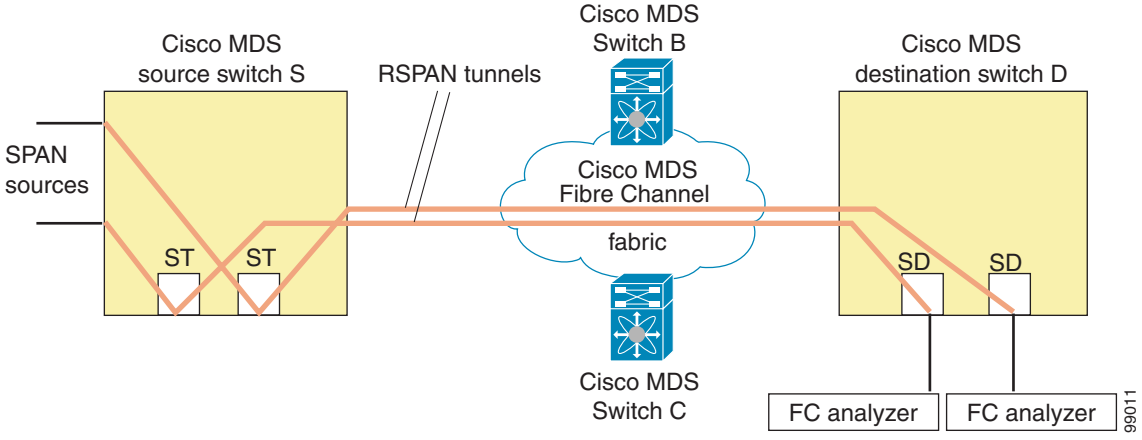

Multiple Sources with Multiple RSPAN Tunnels

Figure 52-18 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 52-18 RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels

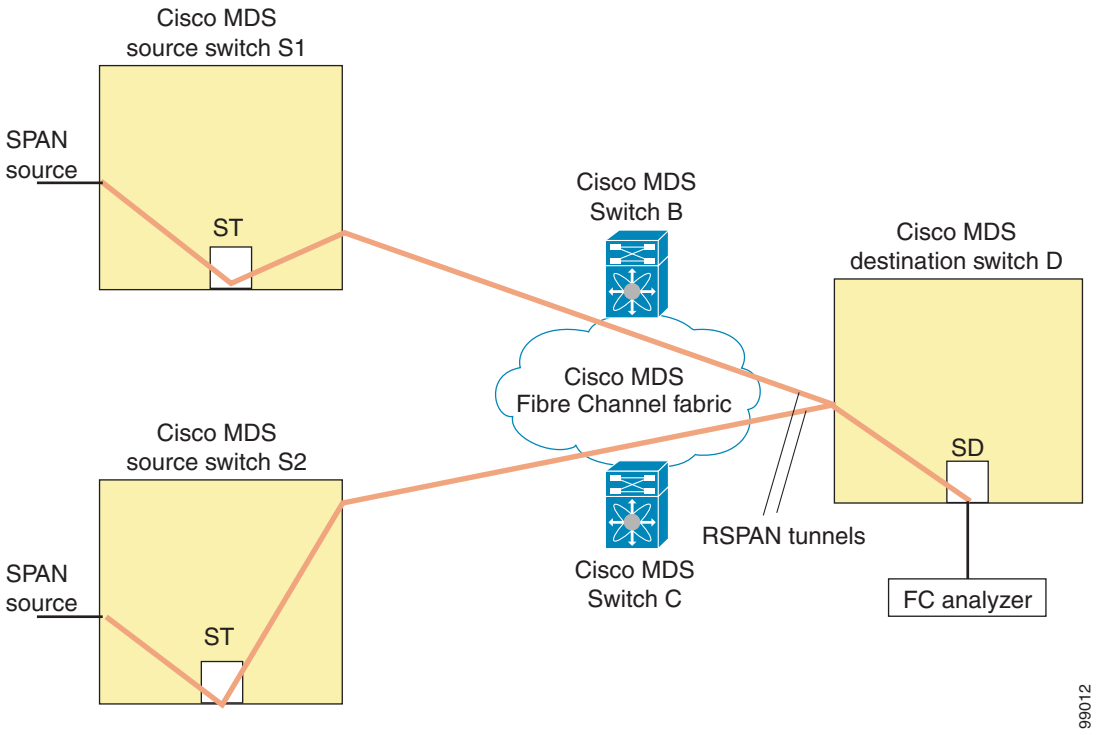

This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Displaying RSPAN Information

Use the **show** commands to display configured RSPAN information. See Examples 52-8 to 52-14.

Example 52-8 Displays ST Port Interface Information

```
switch# show interface brief

--------------------------------------------------------------------------------
Interface   Vsan    Admin    Admin    Status      Oper    Oper    Port-channel
                    Mode     Trunk                Mode    Speed
                             Mode                         (Gbps)
--------------------------------------------------------------------------------
fc1/1       1       auto     on       trunking    TE          2       --
...
fc1/14      1       auto     on       trunking    TE          2       --
fc1/15      1       ST       on       up          ST          2       --
...
```

```
fc2/9       1       auto    on      trunking    TE          2       port-channel 21
fc2/10      1       auto    on      trunking    TE          2       port-channel 21
...
fc2/13      999     auto    on      up          F           1       --
fc2/14      999     auto    on      up          FL          1       --
fc2/15      1       SD      --      up          SD          2       --
fc2/16      1       auto    on      trunking    TE          2       --
--------------------------------------------------------------------------------
Interface           Status      Speed
                                (Gbps)
--------------------------------------------------------------------------------
sup-fc0             up          1

--------------------------------------------------------------------------------
Interface           Status      IP Address              Speed       MTU
--------------------------------------------------------------------------------
mgmt0               up          172.22.36.175/22        100 Mbps    1500

--------------------------------------------------------------------------------
Interface           Status      IP Address              Speed       MTU--
--------------------------------------------------------------------------------
vsan5               up          10.10.10.1/24           1 Gbps      1500
--------------------------------------------------------------------------------
Interface           Vsan        Admin           Status      Oper    Oper
                                Trunk                       Mode    Speed
                                Mode                                (Gbps)
--------------------------------------------------------------------------------
port-channel 21     1           on              trunking    TE      4
--------------------------------------------------------------------------------
Interface           Status      Dest IP Addr    Src IP Addr     TID    Explicit Path
--------------------------------------------------------------------------------
fc-tunnel 100       up          10.10.10.2      10.10.10.1      100
```

Example 52-9 Displays Detailed Information for the ST Port Interface

```
switch# show interface fc1/11
fc1/11 is up
    Hardware is Fibre Channel
    Port WWN is 20:0b:00:05:30:00:59:de
    Admin port mode is ST
    Port mode is ST
    Port vsan is 1
    Speed is 1 Gbps
    Rspan tunnel is fc-tunnel 100
    Beacon is turned off
    5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
    5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
      6862 frames input, 444232 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      6862 frames output, 307072 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

Example 52-10 Displays the FC Tunnel Status

```
switch# show fc-tunnel
fc-tunnel is enabled
```

Example 52-11 Displays FC Tunnel Egress Mapping Information

```
switch# show fc-tunnel tunnel-id-map
tunnel id egress interface
    150     fc3/1
    100     fc3/1
```

Note Multiple tunnel IDs can terminate at the same interface.

Example 52-12 Displays FC Tunnel Explicit Mapping Information

```
switch# show fc-tunnel explicit-path
Explicit path name: Alternate1
       10.20.1.2 loose
       10.20.1.3 strict
Explicit path name: User2
       10.20.50.1 strict
       10.20.50.4 loose
```

Example 52-13 Displays SPAN Mapping Information

```
switch# show span session
Session 2 (active)
   Destination is fc-tunnel 100
   No session filters configured
   Ingress (rx) sources are
     fc2/16,
   Egress (tx) sources are
     fc2/16,
```

Example 52-14 Displays the FC Tunnel Interface

```
switch# show interface fc-tunnel 200
fc-tunnel 200 is up
Dest   IP Addr: 200.200.200.7   Tunnel ID: 200
Source IP Addr: 200.200.200.4   LSP ID: 1
Explicit Path Name:
```

Default SPAN and RSPAN Settings

Table 52-1 lists the default settings for SPAN parameters.

Table 52-1 Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

Table 52-2 lists the default settings for RSPAN parameters.

Table 52-2 Default RSPAN Configuration Parameters

Parameters	Default
FC tunnel	Disabled.
Explicit path	Not configured.
Minimum cost path	Used if explicit path is not configured.

CHAPTER 53

Configuring System Message Logging

This chapter describes how to configure system message logging on Cisco MDS 9000 Family switches. It includes the following sections:

About System Message Logging

properly configured system message logging server. You can also monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a system message logging server.

Note When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a system message logging server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Table 53-1 describes some samples of the facilities supported by the system message logs.

Table 53-1 Internal Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
acl	ACL manager	Cisco MDS 9000 Family specific
all	All facilities	Cisco MDS 9000 Family specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
bootvar	Bootvar	Cisco MDS 9000 Family specific
callhome	Call Home	Cisco MDS 9000 Family specific
cron	Cron or at facility	Standard

Table 53-1 Internal Logging Facilities (continued)

Facility Keyword	Description	Standard or Cisco MDS Specific
daemon	System daemons	Standard
fcc	FCC	Cisco MDS 9000 Family specific
fcdomain	fcdomain	Cisco MDS 9000 Family specific
fcns	Name server	Cisco MDS 9000 Family specific
fcs	FCS	Cisco MDS 9000 Family specific
flogi	FLOGI	Cisco MDS 9000 Family specific
fspf	FSPF	Cisco MDS 9000 Family specific
ftp	File Transfer Protocol	Standard
ipconf	IP configuration	Cisco MDS 9000 Family specific
ipfc	IPFC	Cisco MDS 9000 Family specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard
lpr	Line printer system	Standard
mail	Mail system	Standard
mcast	Multicast	Cisco MDS 9000 Family specific
module	Switching module	Cisco MDS 9000 Family specific
news	USENET news	Standard
ntp	NTP	Cisco MDS 9000 Family specific
platform	Platform manager	Cisco MDS 9000 Family specific
port	Port	Cisco MDS 9000 Family specific
port-channel	PortChannel	Cisco MDS 9000 Family specific
qos	QoS	Cisco MDS 9000 Family specific
rdl	RDL	Cisco MDS 9000 Family specific
rib	RIB	Cisco MDS 9000 Family specific
rscn	RSCN	Cisco MDS 9000 Family specific
securityd	Security	Cisco MDS 9000 Family specific
syslog	Internal system messages	Standard
sysmgr	System manager	Cisco MDS 9000 Family specific
tlport	TL port	Cisco MDS 9000 Family specific
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard
vhbad	Virtual host base adapter daemon	Cisco MDS 9000 Family specific
vni	Virtual network interface	Cisco MDS 9000 Family specific
vrrp_cfg	VRRP configuration	Cisco MDS 9000 Family specific
vrrp_eng	VRRP engine	Cisco MDS 9000 Family specific
vsan	VSAN system messages	Cisco MDS 9000 Family specific

Table 53-1 Internal Logging Facilities (continued)

Facility Keyword	Description	Standard or Cisco MDS Specific
vshd	vshd	Cisco MDS 9000 Family specific
wwn	WWN manager	Cisco MDS 9000 Family specific
xbar	Xbar system messages	Cisco MDS 9000 Family specific
zone	Zone server	Cisco MDS 9000 Family specific

Table 53-2 describes the severity levels supported by the system message logs.

Table 53-2 Error Message Severity Levels

Level Keyword	Level	Description	System Message Definition
emergencies	0	System unusable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

Note Refer to the *Cisco MDS 9000 Family System Messages Reference* for details on the error log message format.

System Message Logging Configuration

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

This sections includes the following topics:

Message Logging Initiation

You can disable logging to the console or enable logging to a given Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet or SSH session, follow these steps:

	Command	Purpose
Step 1	`switch# terminal monitor`	Enables logging for a Telnet or SSH session. **Note** A console session is enabled by default.
Step 2	`switch# terminal no monitor`	Disables logging for a Telnet or SSH session. **Note** A Telnet or SSH session is disabled by default.

Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).

Tip The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud.

See the "Configuring Console Port Settings" section on page 5-28.

To configure the severity level for the console session, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging console 3`	Configures console logging at level 3 (error). Logging messages with a severity level of 3 or above are displayed on the console.
	`switch(config)# no logging console`	Reverts console logging to the factory set default severity level of 2 (critical). Logging messages with a severity level of 2 or above are displayed on the console.

Monitor Severity Level

When logging is enabled for a monitor session (default), you can configure the severity levels of messages that appear on the monitor. The default severity for monitor logging is 5 (notifications).

To configure the severity level for a monitor session, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging monitor 3`	Configures monitor logging at level 3 (error). Logging messages with a severity level of 3 or above are displayed on the monitor.
	`switch(config)# no logging monitor`	Reverts monitor logging to the factory set default severity level of 5 (notifications). Logging messages with a severity level of 5 or above are displayed on the console.

Module Logging

By default, logging is enabled at level 7 for all modules. You can enable or disable logging for each module at a specified level.

To enable or disable the logging for modules and configure the severity level, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging module 1`	Configures module logging at level 1 (alerts) for all modules.
	`switch(config)# logging module`	Configures module logging for all modules in the switch at the default level 5 (notifications).
	`switch(config)# no logging module`	Disables module logging.

Facility Severity Levels

To configure the severity level for a logging facility (see Table 53-1), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging level kernel 4`	Configures Telnet or SSH logging for the kernel facility at level 4 (warning). As a result, logging messages with a severity level of 4 or above are displayed.
	`switch(config)# no logging level kernel 4`	Reverts to the default severity level 6 (informational) for the Telnet or SSH logging for the kernel facility. **Note** Use the **show logging info** command to display the default logging levels for the facilities listed in Table 53-1.

Log Files

Logging messages can be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages. The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes.

To send log messages to a file, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# logging logfile messages 3`	Configures logging of information for errors or events above with a severity level 3 or above to the default log file named messages.
	`switch(config)# logging logfile ManagerLog 3`	Configures logging of information for errors or events with a severity level 3 or above to a file named ManagerLog using the default size of 10,485,760 bytes.
	`switch(config)# logging logfile ManagerLog 3 size 3000000`	Configures logging information for errors or events with a severity level 3 or above to a file named ManagerLog. By configuring a size, you are restricting the file size to 3,000,000 bytes.
	`switch(config)# no logging logfile`	Disables logging messages to the logfile.

Note You can rename the log file using the **logging logfile** command.

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed. You can use the **show logging logfile** and **clear logging logfile** commands to view and delete the contents of this file. You can use the **dir log:** command to view logging file statistics. You can use the **delete log:** command to remove the log file.

You can copy the logfile to a different location using the **copy log:** command using additional copy syntax (see the "Copying Configuration Files" section on page 8-5).

System Message Logging Servers

You can configure a maximum of three system message logging servers.

To send log messages to a UNIX system message logging server, you must configure the system message logging daemon on a UNIX server. Log in as root, and follow these steps:

Step 1 Add the following line to the /etc/syslog.conf file.

local1.debug **/var/log/*myfile*.log**

Note Be sure to add five tab characters between **local1.debug** and **/var/log/*myfile*.log**. Refer to entries in the /etc/syslog.conf file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```

Step 3 Make sure the system message logging daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

To configure system message logging server IPv4 addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch#`	Enters configuration mode.
Step 2	`switch(config)# logging server 172.22.00.00`	Configures the switch to forward log messages according to the specified facility types and severity levels to remote multiple servers specified by its hostname or IPv4 address (172.22.00.00).
	`switch(config)# logging server 172.22.00.00 facility local1`	Configures the switch to forward log messages according to the specified facility (local1) for the server IPv4 address (172.22.00.00). The default outgoing facility is local7.
	`switch(config)# no logging server 172.11.00.00`	Removes the specified server (172.11.00.00) and reverts to factory default.

To configure system message logging server IPv6 addresses, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch#`	Enters configuration mode.
Step 2	`switch(config)# logging server 2001::0db8:800:200c:417a`	Configures the switch to forward log messages according to the specified facility types and severity levels to a remote server specified by its IPv6 address.
	`switch(config)# logging server 2001::0db8:800:200c:417a facility local1`	Configures the switch to forward log messages according to the specified facility (local1) for the server IPv6 address. The default outgoing facility is local7.
	`switch(config)# no logging server 2001::0db8:800:200c:417a`	Removes the specified server and reverts to factory default.

Outgoing System Message Logging Server Facilities

All system messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single system message logging daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in Table 53-1 and the outgoing logging facilities are listed in Table 53-3.

Table 53-3 Outgoing Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
cron	Cron or at facility	Standard
daemon	System daemons	Standard
ftp	File Transfer Protocol	Standard
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal system messages	Standard
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard

System Message Logging Configuration Distribution

You can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform system message logging configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The system message logging server uses the effective and pending database model to store or commit the commands based on your configuration. When you commit the 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. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. See Chapter 6, “Using the CFS Infrastructure” for more information on the CFS application.

To enable fabric distribution for system message logging server configurations, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# logging distribute`	Enables the system message logging server configuration to be distributed to all switches in the fabric, acquires a lock, and stores all future configuration changes in the pending database.
	`switch(config)# no logging distribute`	Disables (default) system message logging server configuration distribution to all switches in the fabric.

To commit the system message logging server configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# logging commit`	Distributes the configuration changes to all switches in the fabric, releases the lock, and overwrites the effective database with the changes made to the pending database.

To discard the system message logging server configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# logging abort`	Discards the system message logging server configuration changes in the pending database and releases the fabric lock.

Fabric Lock Override

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

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

To use administrative privileges and release a locked system message logging session, use the **clear logging session** command.

```
switch# clear logging session
```

Database Merge Guidelines

See the "CFS Merge Support" section on page 6-8 for detailed concepts.

When merging two system message logging databases, follow these guidelines:

- Be aware that the merged database is a union of the existing and received database for each switch in the fabric.
- Verify that the merged database will only have a maximum of three system message logging servers.

Caution If the merged database contains more that three servers, the merge will fail.

Displaying System Message Logging Information

Use the **show logging** command to display the current system message logging configuration. See Examples 53-1 to 53-10.

Note When using the **show logging** command, output is displayed only when the configured logging levels for the switch are different from the default levels.

Example 53-1 Displays Current System Message Logging

```
switch# show logging
Logging console:                 enabled (Severity: critical)
Logging monitor:                 enabled (Severity: debugging)
Logging linecard:                enabled (Severity: debugging)
Logging server:                  enabled
{172.20.102.34}
        server severity:         debugging
        server facility:         local7
{10.77.202.88}
        server severity:         debugging
        server facility:         local7
{10.77.202.149}
        server severity:         debugging
        server facility:         local7
Logging logfile:                 enabled
        Name - messages: Severity - debugging Size - 4194304
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
```

```
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)

Feb 14 09:50:57 excal-113 %TTYD-6-TTYD_MISC: TTYD TTYD started
Feb 14 09:50:58 excal-113 %DAEMON-6-SYSTEM_MSG: precision = 8 usec
...
```

Use the **show logging nvram** command to view the log messages saved in NVRAM. Only log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Example 53-2 Displays NVRM Log Contents

```
switch# show logging nvram
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2209, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2199, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
...
```

Example 53-3 Displays the Log File

```
switch# show logging logfile
Jul 16 21:06:50 %DAEMON-3-SYSTEM_MSG: Un-parsable frequency in /mnt/pss/ntp.drift
Jul 16 21:06:56 %DAEMON-3-SYSTEM_MSG: snmpd:snmp_open_debug_cfg: no snmp_saved_dbg_uri ;
Jul 16 21:06:58 172.22.91.204 %PORT-5-IF_UP: Interface mgmt0 is up
Jul 16 21:06:58 172.22.91.204 %MODULE-5-ACTIVE_SUP_OK: Supervisor 5 is active
...
```

Example 53-4 Displays Console Logging Status

```
switch# show logging console
Logging console:                enabled (Severity: notifications)
```

Example 53-5 Displays Logging Facility

```
switch# show logging level
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
```

```
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Example 53-6 Displays Logging Information

```
switch# show logging info
Logging console:                enabled (Severity: critical)
Logging monitor:                enabled (Severity: debugging)
Logging linecard:               enabled (Severity: debugging)
Logging server:                 enabled
{172.20.102.34}
        server severity:        debugging
        server facility:        local7
{10.77.202.88}
        server severity:        debugging
        server facility:        local7
{10.77.202.149}
        server severity:        debugging
        server facility:        local7
Logging logfile:                enabled
        Name - messages: Severity - debugging Size - 4194304
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
```

```
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Example 53-7 Displays Last Few Lines of a Log File

```
switch# show logging last 2
Nov 8 16:48:04 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/1
(171.71.58.56)
Nov 8 17:44:09 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(171.71.58.72)
```

Example 53-8 Displays Switching Module Logging Status

```
switch# show logging module
Logging linecard:                enabled (Severity: debugging)
```

Example 53-9 Displays Monitor Logging Status

```
switch# show logging monitor
Logging monitor:                 enabled (Severity: information)
```

Example 53-10 Displays Server Information

```
switch# show logging server
Logging server:                  enabled
{172.22.95.167}
        server severity:         debugging
        server facility:         local7
{172.22.92.58}
        server severity:         debugging
```

```
server facility:        local7
```

Default Settings

Table 53-4 lists the default settings for system message logging.

Table 53-4 Default System Message Log Settings

Parameters	Default
System message logging to the console	Enabled for messages at the critical severity level.
System message logging to Telnet sessions	Disabled.
Logging file size	4194304.
Log file name	Message (change to a name with up to 200 characters).
Logging server	Disabled.
Syslog server IP address	Not configured.
Number of servers	Three servers.
Server facility	Local 7.

CHAPTER 54

Configuring Call Home

Call Home provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, e-mail notification to a Network Operations Center, and utilization of Cisco AutoNotify services for direct case generation with the Technical Assistance Center.

The Call Home feature provides message throttling capabilities. Periodic inventory messages, port syslog messages, and RMON alert messages are added to the list of deliverable Call Home messages. If required you can also use the Cisco Fabric Services application to distribute the Call Home configuration to all other switches in the fabric.

This chapter includes the following sections:

Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family. It provides multiple Call Home profiles (also referred to as *Call Home destination profiles*), each with separate potential destinations. You can define your own destination profiles in addition to predefined profiles.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
  - Short Text—Suitable for pagers or printed reports.
  - Plain Text—Full formatted message information suitable for human reading.
  - XML—Matching readable format using Extensible Markup Language (XML) and document type definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco.com website at http://www.cisco.com/. The XML format enables communication with the Cisco Systems Technical Assistance Center.
- Multiple concurrent message destinations. You can configure up to 50 e-mail destination addresses for each destination profile.
- Multiple message categories including system, environment, switching module hardware, supervisor module, hardware, inventory, syslog, RMON, and test.

Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible by registering with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support.

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, e-mail server, and an XML destination profile as specified in the Service Activation document found on the Cisco.com web site at:
http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccmsrvs/sssrvact.htm

To configure a Cisco MDS 9000 Family switch to use the AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and e-mail address information is found on the Cisco.com web site at:
http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_configuration_example09186a0080108e72.shtml

To register, the following items are required:

- The SMARTnet contract number covering your Cisco MDS 9000 Family switch.
- Your name, company address, your e-mail address, and your Cisco.com ID.
- The exact product number of your Cisco MDS 9000 Family switch. For example, valid product numbers include DS-C6509 and DS-C9216-K9.
- The serial number of your Cisco MDS 9000 Family switch. This can be obtained by looking at the serial number label on the back of the switch (next to the power supply).

The ContractID, CustomerID, SiteID, and SwitchPriority parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

Use the **show sprom backplane 1** command or the s**how license host-id** command to obtain the switch serial number.

Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- An e-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, e-mail, or automated service such as Cisco AutoNotify.
- Switches can forward events (SNMP traps/informs) up to 10 destinations.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This is required to determine the origin of messages received.
- The Cisco MDS 9000 switch must have IP connectivity to an e-mail server.
- If Cisco AutoNotify is used, an active service contract must cover the device being configured.

To configure Call Home, follow these steps:

Step 1 Assign contact information.

Step 2 Configure destination profiles.

Step 3 Associate one or more alert groups to each profile as required by your network. Customize the alert groups, if desired.

Step 4 Configure e-mail options.

Step 5 Enable or disable Call Home.

Step 6 Test Call Home messages.

Contact Information

It is mandatory for each switch to include e-mail, phone, and street address information. It is optional to include the contract ID, customer ID, site ID, and switch priority information.

To assign the contact information, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch# snmp-server contact personname@companyname.com`	Configures the SNMP contact name.
Step 3	`switch(config)# callhome` `switch(config-callhome)#`	Enters the Call Home configuration submode.
Step 4	`switch(config-callhome)# email-contact username@company.com`	Assigns the customer's e-mail address. Up to 128 alphanumeric characters are accepted in e-mail address format. **Note** You can use any valid e-mail address. You cannot use spaces.
Step 5	`switch(config-callhome)# phone-contact +1-800-123-4567`	Assigns the customer's phone number. Up to 20 alphanumeric characters are accepted in international format. **Note** You cannot use spaces. Be sure to use the + prefix before the number.
Step 6	`switch(config-callhome)# streetaddress 1234 Picaboo Street, Any city, Any state, 12345`	Assigns the customer's street address where the equipment is located. Up to 256 alphanumeric characters are accepted in free format.
Step 7	`switch(config-callhome)# switch-priority 0`	Assigns the switch priority, with 0 being the highest priority and 7 the lowest. **Tip** Use this field to create a hierarchical management structure.
Step 8	`switch(config-callhome)# customer-id Customer1234`	Optional. Identifies the customer ID. Up to 256 alphanumeric characters are accepted in free format.
Step 9	`switch(config-callhome)# site-id Site1ManhattanNY`	Optional. Identifies the customer site ID. Up to 256 alphanumeric characters are accepted in free format.
Step 10	`switch(config-callhome)# contract-id Company1234`	Assigns the customer ID for the switch. Up to 64 alphanumeric characters are accepted in free format.

Note Switches can forward events (SNMP traps/informs) up to 10 destinations.

Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. At least one destination profile is required. You can configure multiple destination profiles of one or more types.

You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.

Note If you use the Cisco AutoNotify service, the XML destination profile is required (see http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_configuration_example09186a0080108e72.shtml).

You can configure the following attributes for a destination profile:

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).

To configure predefined destination profile messaging options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters the Call Home configuration submode.
Step 3	`switch(config-callhome)#` **`destination-profile full-txt-destination email-addr person@place.com`**	Configures an e-mail address for the predefined full-txt-destination profile. The e-mail addresses in this destination profile receives messages in full-txt format. The full-text format provides the complete, detailed explanation of the failure. **Tip** Use a standard e-mail address that does not have any text size restrictions.
	`switch(config-callhome)#` **`destination-profile full-txt-destination message-size 1000000`**	Configures a maximum destination message size for the predefined full-txt-destination profile. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.
Step 4	`switch(config-callhome)#` **`destination-profile short-txt-destination email-addr person@place.com`**	Configures an e-mail address for the predefined short-txt-destination profile. The e-mail addresses in this destination profile receive messages in short-txt format. This format provides the basic explanation of the failure in the Call Home message. **Tip** Use a pager-related e-mail address for this option.
	`switch(config-callhome)#` **`destination-profile short-txt-destination message-size 100000`**	Configures maximum destination message size for the predefined short-txt-destination profile. The valid range is 0 to 1,000,000 bytes and the default is 4000. A value of 0 implies that a message of any size can be sent.

	Command	Purpose
Step 5	`switch(config-callhome)# destination-profile XML-destination email-addr findout@.cisco.com`	Configures an e-mail address for the predefined XML-destination profile. The e-mail addresses in this destination-profile receives messages in XML format. This format provides information that is compatible with Cisco Systems TAC support. **Tip** Do not add a pager-related e-mail address to this destination profile because of the large message size.
	`switch(config-callhome)# destination-profile XML-destination message-size 100000`	Configures maximum destination message size for the predefined destination profile `XML-destination`. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.

Note Steps 3, 4, and 5 in this procedure can be skipped or configured in any order.

To configure a new destination-profile (and related parameters), follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters the Call Home configuration submode.
Step 3	`switch(config-callhome)# destination-profile test`	Configures a new destination profile called test.
Step 4	`switch(config-callhome)# destination-profile test email-addr person@place.com`	Configures the e-mail address for the user-defined destination profile (test) sent in default XML format.
Step 5	`switch(config-callhome)# destination-profile test message-size 1000000`	Configures a maximum message size for the destination e-mail addresses in the user-defined destination profile (test) sent in default XML format. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.
Step 6	`switch(config-callhome)# destination-profile test format full-txt`	Configures message-format for the user-defined destination profile (test) to be full text format.
	`switch(config-callhome)# destination-profile test format short-txt`	Configures message-format for the user-defined destination profile (test) to be short text format.

Note Steps 4, 5, and 6 in this procedure can be skipped or configured in any order.

Alert Groups

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family. Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

The alert group feature allows you to select the set of Call Home alerts to be received by a destination profile (either predefined or user-defined). You can associate multiple alert groups with a destination profile.

Note A Call Home alert is sent to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

To associate an alert group with a destination profile, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# destination-profile test1 alert-group test`	Optional. Configures user-defined destination profile (test1) to receive all user-generated Call Home test notifications.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group test`	Optional. Configures predefined short-text destination profile to receive all user-generated Call Home test notifications.
Step 4	`switch(config-callhome)# destination-profile test1 alert-group all`	Optional. Configures user-defined destination profile (test1) to receive Call Home notifications for all events
	`switch(config-callhome)# destination-profile short-txt-destination alert-group all`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for all (default) events
Step 5	`switch(config-callhome)# destination-profile test1 alert-group Cisco-TAC`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for events that are meant only for Cisco TAC or the Auto-notify service.
	`switch(config-callhome)# destination-profile xml-destination alert-group Cisco-TAC`	Optional. Configures predefined XML destination message profile to receive Call Home notifications for events that are meant only for Cisco TAC or the auto-notify service.
Step 6	`switch(config-callhome)# destination-profile test1 alert-group environmental`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for power, fan, and temperature-related events.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group environmental`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for power, fan, and temperature-related events.

	Command	Purpose
Step 7	`switch(config-callhome)# destination-profile test1 alert-group inventory`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for inventory status events.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group inventory`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for inventory status events.
Step 8	`switch(config-callhome)# destination-profile test1 alert-group linecard-hardware`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for module-related events.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group linecard-hardware`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for module-related events.
Step 9	`switch(config-callhome)# destination-profile test1 alert-group supervisor-hardware`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for supervisor-related events.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group supervisor-hardware`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for supervisor-related events.
Step 10	`switch(config-callhome)# destination-profile test1 alert-group system`	Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for software-related events.
	`switch(config-callhome)# destination-profile short-txt-destination alert-group system`	Optional. Configures predefined short-text destination message profile to receive Call Home notifications for software-related events.

Customized Alert Group Messages

The predefined Call Home alert groups generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional valid **show** commands when specific events occur. The output from these additional **show** commands is included in the notification message along with that of the predefined **show** commands.

Note You can assign a maximum of five user-defined **show** commands to an alert group. Only **show** commands can be assigned to an alert group.

Note Customized show commands are only supported for full text and XML alert groups. Short text alert groups (short-txt-destination) do not support customized **show** commands because they only allow 128 bytes of text.

To assign **show** commands to be executed when an alert is sent, you must associate the commands with the alert group. When an alert is sent, Call Home associates the alert group with an alert type and attaches the output of the **show** commands to the alert message.

Note Make sure the destination profiles for a non-Cisco-TAC alert group, with a predefined **show** command, and the Cisco-TAC alert group are not the same.

To customize Call Home alert group messages, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# alert-group license user-def-cmd "show license usage"`	Configures a user-defined **show** command for an alert group license. **Note** The **show** command must be enclosed in double quotes. Only valid **show** commands are accepted.
	`switch(config-callhome)# no alert-group license user-def-cmd "show license usage"`	Removes the user-defined **show** command from the alert group.

Verifying Alert Group Customization

To verify the alert group customization, use the **show callhome user-def-cmds** command.

```
switch# show callhome user-def-cmds
User configured commands for alert groups :
alert-group test user-def-cmd "show version"
```

Call Home Message Level Feature

The Call Home message level feature allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower that the urgency threshold is not sent. The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages are sent).

Note Call Home severity levels are not the same as system message logging severity levels.

To set the message level for each destination profile for Call Home, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.

	Command	Purpose
Step 3	`switch(config-callhome)# destination-profile test message-level 5`	Optional. Configures the message level urgency as 5 and above for the user-defined profile (test1).
	`switch(config-callhome)# no destination-profile oldtest message-level 7`	Removes a previously configured urgency level and reverts it to the default of 0 (all messages are sent).

Syslog-Based Alerts

You can configure the switch to send certain syslog messages as Call Home messages. The syslog-group-port alert group selects syslog messages for the port facility. The Call Home application maps the syslog severity level to the corresponding Call Home severity level (see the "Call Home Message Levels" section on page 54-22). For example, if you select level 5 for the Call Home message level, syslog messages at levels 0, 1, and 2 are included in the Call Home log.

Whenever a syslog message is generated, the Call Home application sends a Call Home message depending on the mapping between the destination profile and the alert group mapping and based on the severity level of the generated syslog message. To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level (see the "Call Home Message Level Feature" section on page 54-9).

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Family System Messages Guide.*

To configure the syslog-group-port alert group, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# destination-profile short-txt-destination alert-group syslog-group-port`	Configures the predefined destination profile (short-txt-destination) to receive Call Home Notifications corresponding to syslog messages for the port facility.
Step 4	`switch(config-callhome)# destination-profile short-txt-destination message-level 5`	Optional. Configures the predefined destination-profile (short-txt-destination) to send a Call Home message for syslog messages whose severity levels map to Call Home severity level of 5 or greater. The default is message level 0 (all syslog messages).

RMON-Based Alerts

You can configure the switch to send Call Home notifications corresponding to RMON alert triggers. All RMON-based Call Home messages have their message level set to NOTIFY (2). The RMON alert group is defined for all RMON-based Call Home alerts. To receive an RMON-based Call Home alert, you must associate a destination profile with the RMON alert group.

To configure RMON alert groups, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# destination-profile xml-destination alert-group rmon`	Optional. Configures a destination message profile (rmon_group) to send Call Home notifications for configured RMON messages.

E-Mail Options

You can configure the from, reply-to, and return-receipt e-mail addresses. While most e-mail address configurations are optional, you must configure the SMTP server address for the Call Home functionality to work.

Configuring General E-Mail Options

To configure general e-mail options, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# transport email from user@company1.com`	Optional. Configures the from e-mail address.
Step 4	`switch(config-callhome)# transport email reply-to person@place.com`	Optional. Configures the reply-to e-mail address to which all responses should be sent.

Configuring SMTP Server and Ports

To configure the SMTP server and port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.

	Command	Purpose
Step 3	`switch(config-callhome)# transport email smtp-server 192.168.1.1`	Configures the DNS, IPv4 address, or IPv6 address of the SMTP server to reach the server. The port usage defaults to 25 if no port is specified. **Note** The port number is optional and, if required, may be changed depending on the server location.
	`switch(config-callhome)# transport email smtp-server 192.168.1.1 port 30`	

Periodic Inventory Notification

You can configure the switch to periodically send a message with an inventory of all the software services currently enabled and running on the switch along with hardware inventory information. The inventory is modified each time the switch is restarted nondisruptively.

By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. When you enable this feature without configuring an interval value, the Call Home message is sent every 7 days. This value ranges from 1 to 30 days.

To enable periodic inventory notification in a Cisco MDS 9000 Family switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters the Call Home configuration submode.
Step 3	`switch(config-callhome)# periodic-inventory notification`	Enables the periodic inventory notification feature. By default, the Call Home message is sent every 7 days.
	`switch(config-callhome)# no periodic-inventory notification`	Disables the periodic inventory notification feature (default).
Step 4	`switch(config-callhome)# periodic-inventory notification interval 15`	Configures the periodic inventory notification message to be sent every 15 days. This value ranges from 1 to 30 days.
	`switch(config-callhome)# no periodic-inventory notification interval 15`	Defaults to using the factory default of sending a Call Home message every 7 days.

Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family. When enabled, if the number of messages sent exceeds the maximum limit of 30 messages within the 2-hour time frame, then further messages for that alert type are discarded within that time frame. You cannot modify the time frame or the message counter limit.

If 2 hours have elapsed since the first such message was sent and a new message has to be sent, then the new message is sent and the time frame is reset to the time when the new message was sent and the count is reset to 1.

To enable message throttling in a Cisco MDS 9000 Family switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters the Call Home configuration submode.
Step 3	`switch(config-callhome)# no duplicate-message throttle`	Disables the duplicate message throttling feature.
	`switch(config-callhome)# duplicate-message throttle`	Enables the duplicate message throttling feature (default).

Call Home Enable Function

Once you have configured the contact information, you must enable the Call Home function.

To enable the Call Home function, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.

Step 3	`switch(config-callhome)# enable`
	`callhome enabled successfully`
	`switch(config-callhome)#`
		Enables the Call Home function.

	`switch(config-callhome)# disable` `switch(config-callhome)#`	Disables the Call Home function. When you disable the Call Home function, all input events are ignored. **Note** Even if Call Home is disabled, basic information for each Call Home event is sent.

Call Home Configuration Distribution

You can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform Call Home configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The Call Home application uses the effective and pending database model to store or commit the configuration changes. When you commit the 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. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. See Chapter 6, "Using the CFS Infrastructure," for more information on the CFS application.

Note The Switch priority and the Syscontact name are not distributed.

To enable Call Home fabric distribution, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# distribute`	Enables Call Home configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.
	`switch(config-callhome)# no distribute`	Disables (default) Call Home configuration distribution to all switches in the fabric.

To commit the Call Home configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# commit`	Distributes the configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database.

To discard the Call Home configuration changes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# callhome` `switch(config-callhome)#`	Enters Call Home configuration submode.
Step 3	`switch(config-callhome)# abort`	Discards the configuration changes in the pending database and releases the fabric lock.

Fabric Lock Override

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

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

To use administrative privileges and release a locked Call Home session, use the **clear callhome session** command.

```
switch# clear callhome session
```

Database Merge Guidelines

See the "CFS Merge Support" section on page 6-8 for detailed concepts.

When merging two Call Home databases, follow these guidelines:

- Be aware that the merged database contains the following information:
  - A superset of all the destination profiles from the dominant and subordinate switches take part in the merge protocol.
  - The e-mail addresses and alert groups for the destination profiles.
  - Other configuration information (for example, message throttling, periodic inventory) from the switch that existed in the dominant switch before the merge.
- Verify that two destination profiles do not have the same name (even if they have different configuration information) on the subordinate and dominant switches. If they do contain the same name, the merge operation will fail. You must then modify or delete the conflicting destination profile on the required switch.

Call Home Communications Test

Use the **test** command to simulate a message generation.

To test the Call Home function, follow these steps:

	Command	Purpose

Step 1	`switch# callhome test`
	`trying to send test callhome message`
	`successfully sent test callhome message`
		Sends a test message to the configured destination(s).

Step 2	`switch# callhome test inventory`
	`trying to send test callhome message`
	`successfully sent test callhome message`
		Sends a test inventory message to the configured destination(s).

Displaying Call Home Information

Use the **show callhome** command to display the configured Call Home information (see Examples 54-1 to 54-7).

Example 54-1 Displays Configured Call Home Information

```
switch# show callhome
callhome enabled
Callhome Information:
contact person name:who@where
contact person's email:person@place.com
contact person's phone number:310-408-4000
street addr:1234 Picaboo Street, Any city, Any state, 12345
site id:Site1ManhattanNewYork
customer id:Customer1234
contract id:Cisco1234
switch priority:0
```

Example 54-2 Displays Information for All Destination Profiles (Predefined and User-Defined)

```
switch# show callhome destination-profile
XML destination profile information
maximum message size:500000
message format:XML
message-level:0
email addresses configured:
alert groups configured:
cisco_tac

test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
admin@yourcompany.com

alert groups configured:
test

full-txt destination profile information
maximum message size:500000
message format:full-txt
message-level:0
email addresses configured:

alert groups configured:
all

short-txt destination profile information
maximum message size:4000
message format:short-txt
message-level:0
email addresses configured:

alert groups configured:
all
```

Example 54-3 Displays Information for a User-defined Destination Profile

```
switch# show callhome destination-profile test
test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
user@company.com

alert groups configured:
test
```

Example 54-4 Displays the Full-Text Profile

```
switch# show callhome destination-profile profile full-txt-destination
full-txt destination profile information
maximum message size:250000
email addresses configured:
person2@company2.com
```

Example 54-5 Displays the Short-Text Profile

```
switch# show callhome destination-profile profile short-txt-destination
Short-txt destination profile information
maximum message size:4000
email addresses configured:
person2@company2.com
```

Example 54-6 Displays the XML Destination Profile

```
switch# show callhome destination-profile profile XML-destination
XML destination profile information
maximum message size:250000
email addresses configured:
findout@.cisco.com
```

Example 54-7 Displays E-Mail and SMTP Information

```
switch# show callhome transport-email
from email addr:user@company1.com
reply to email addr:pointer@company.com
return receipt email addr:user@company1.com
smtp server:server.company.com
smtp server port:25
```

Note Switches can forward events (SNMP traps/informs) up to 10 destinations.

Sample Syslog Alert Notification in Full-txt Format

```
source:MDS9000
Switch Priority:7
Device Id:DS-C9506@C@FG@07120011
```

```
Customer Id:basu
Contract Id:123
Site Id:San Jose
Server Id:DS-C9506@C@FG@07120011
Time of Event:2004-10-08T11:10:44
Message Name:SYSLOG_ALERT
Message Type:Syslog
Severity Level:2
System Name:10.76.100.177
Contact Name:Basavaraj B
Contact Email:admin@yourcompany.com
Contact Phone:+91-80-310-1718
Street Address:#71 , Miller's Road
Event Description:2004 Oct  8 11:10:44 10.76.100.177 %PORT-5-IF_TRUNK_UP: %$VSAN 1%$
Interface fc2/5, vsan 1 is up

syslog_facility:PORT
start chassis information:
Affected Chassis:DS-C9506
Affected Chassis Serial Number:FG@07120011
Affected Chassis Hardware Version:0.104
Affected Chassis Software Version:3.1(1)
Affected Chassis Part No:73-8607-01
end chassis information:
```

Sample Syslog Alert Notification in XML Format

```
X-Mozilla-Status2: 02000000
Return-Path: <tester@cisco.com>
...

<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:SYSLOG_ALERT
-->
<mml>
<header>
<time>2004-09-30T06:12:36</time>
<name>SYSLOG_ALERT</name>
<type>Syslog</type>
<level>2</level>
<source>MDS9000</source>
<priority>7</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>911</custId>
<contractId>33445</contractId>
<siteId>91111</siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>2004 Sep 30 06:12:36 switch186 %PORT-5-IF_UP: %$VSAN 2000%$ Interface fc1/10 is
up in mode FL
</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>admin@yourcompany.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-8888888</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
```

```
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>3.1(1)</swVersion>
</chassis>
<nvp>
<name>syslog_facility</name>
<value>PORT</value>
</nvp>
</body>
</mml>
```

Sample RMON Notification in XML Format

```
Return-Path: <tester@cisco.com>
...
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:RMON_ALERT
-->
<mml>
<header>
<time>2004-10-12T04:59:13</time>
<name>RMON_ALERT</name>
<type>RMON</type>
<level>2</level>
<source>MDS9000</source>
<priority>3</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>0</custId>
<contractId>u</contractId>
<siteId>&amp;</siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>rlaxmina-w2k07</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>admin@yourcompany.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-000000</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>3.1(1)</swVersion>
</chassis>
<nvp>
<name>ThresholdType</name>
<value>RisingThreshold</value>
</nvp>
<nvp>
<name>ThresholdValue</name>
<value>0</value>
</nvp>
<nvp>
<name>AlarmValue</name>
<value>0</value>
</nvp>
```

```
</body>
</mml>
```

Default Settings

Table 54-1 lists the default Call Home settings.

Table 54-1 Default Call Home Settings

Parameters	Default
Destination message size for a message sent in full text format.	500,000.
Destination message size for a message sent in XML format.	500,000.
Destination message size for a message sent in short text format.	4000.
DNS or IP address of the SMTP server to reach the server if no port is specified.	25.
Alert group association with profile.	All.
Format type.	XML.
Call Home message level.	0 (zero).

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned CLI commands to execute when the event occurs. The command output is included in the transmitted message. Table 54-2 lists the trigger events.

Table 54-2 Event Triggers

Event	Alert Group	Event Name	Description	Call Home Message Level
Call Home	System and CISCO_TAC	SW_CRASH	A software process has crashed with a stateless restart, indicating an interruption of a service.	5
	System and CISCO_TAC	SW_SYSTEM_INCONSISTENT	Inconsistency detected in software or file system.	5
	Environmental and CISCO_TAC	TEMPERATURE_ALARM	Thermal sensor indicates temperature reached operating threshold.	6
		POWER_SUPPLY_FAILURE	Power supply failed.	6
		FAN_FAILURE	Cooling fan has failed.	5
	Switching module and CISCO_TAC	LINECARD_FAILURE	Switching module operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Switching module failed power-up diagnostics.	7
	Line Card Hardware and CISCO_TAC	PORT_FAILURE	Hardware failure of interface port(s).	6
	Line Card Hardware, Supervisor Hardware, and CISCO_TAC	BOOTFLASH_FAILURE	Failure of boot compact Flash card.	6
	Supervisor module and CISCO_TAC	SUP_FAILURE	Supervisor module operation failed.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Supervisor module failed power-up diagnostics.	7
	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of in-band communications path.	7
	Supervisor Hardware and CISCO_TAC	EOBC_FAILURE	Ethernet out-of-band channel communications failure.	6
	Supervisor Hardware and CISCO_TAC	MGMT_PORT_FAILURE	Hardware failure of management Ethernet port.	5
	License	LICENSE_VIOLATION	Feature in use is not licensed, and are turned off after grace period expiration.	6

Table 54-2 Event Triggers (continued)

Event	Alert Group	Event Name	Description	Call Home Message Level
Inventory	Inventory and CISCO_TAC	COLD_BOOT	Switch is powered up and reset to a cold boot sequence.	2
		HARDWARE_INSERTION	New piece of hardware inserted into the chassis.	2
		HARDWARE_REMOVAL	Hardware removed from the chassis.	2
Test	Test and CISCO_TAC	TEST	User generated test.	2
Port syslog	Syslog-group-port	SYSLOG_ALERT	Syslog messages corresponding to the port facility.	2
RMON	RMON	RMON_ALERT	RMON alert trigger messages.	2

Table 54-3 lists event categories and command outputs.

Table 54-3 Event Categories and Executed Commands

Event Category	Description	Executed Commands
System	Events generated by failure of a software system that is critical to unit operation.	**show tech-support** **show system redundancy status**
Environmental	Events related to power, fan, and environment sensing elements such as temperature alarms.	**show module** **show environment**
Switching module hardware	Events related to standard or intelligent switching modules.	**show tech-support**
Supervisor hardware	Events related to supervisor modules.	**show tech-support**
Inventory	Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement.	**show version**
Test	User generated test message.	**show version**

Call Home Message Levels

Call Home messages (sent for syslog alert groups) have the syslog severity level mapped to the Call Home message level (see the "Syslog-Based Alerts" section on page 54-10).

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Call Home message levels are preassigned per event type.

Severity levels range from 0 to 9, with 9 having the highest urgency. Each syslog level has keywords and a corresponding syslog level as listed in Table 54-4.

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Family System Messages Guide.*

Note Call Home severity levels are not the same as system message logging severity levels (see Chapter 53, "Configuring System Message Logging" and the *Cisco MDS 9000 Family System Messages Guide*).

Table 54-4 Severity and Syslog Level Mapping

Call Home Level	Keyword Used	Syslog Level	Description
Catastrophic (9)	**Catastrophic**	N/A	Network wide catastrophic failure.
Disaster (8)	**Disaster**	N/A	Significant network impact.
Fatal (7)	**Fatal**	Emergency (0)	System is unusable.
Critical (6)	**Critical**	Alert (1)	Critical conditions, immediate attention needed.
Major (5)	**Major**	Critical (2)	Major conditions.
Minor (4)	**Minor**	Error (3)	Minor conditions.
Warning (3)	**Warning**	Warning (4)	Warning conditions.
Notify (2)	**Notification**	Notice (5)	Basic notification and informational messages. Possibly independently insignificant.
Normal (1)	**Normal**	Information (6)	Normal event signifying return to normal state.
Debug (0)	**Debugging**	Debug (7)	Debugging messages.

Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

Table 54-5 describes the short text formatting option for all message types.

Table 54-5 Short Text Messages

Data Item	Description
Device identification	Configured device name
Date/time stamp	Time stamp of the triggering event
Error isolation message	Plain English description of triggering event
Alarm urgency level	Error level such as that applied to system message

Table 54-6, Table 54-7, and Table 54-8 display the information contained in plain text and XML messages.

Table 54-6 Reactive Event Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DD*T*HH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specific event names are listed in the "Event Triggers" section on page 54-21.	/mml/header/name
Message type	Specifically "Call Home."	/mml/header/type
Message group	Specifically "reactive."	/mml/header/group
Severity level	Severity level of message (see Table 54-4).	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type*@*Sid*@*seria*l, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header/deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header/customerID
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header/siteId
Server ID	If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch. Format: *type*@*Sid*@*seria*l, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event. This is the host name of the device.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail

Table 54-6 Reactive Event Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
Affected FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
Affected FRU serial number	Serial number of affected FRU.	/mml/body/fru/serialNo
Affected FRU part number	Part number of affected FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU generating the event message.	/mml/body/fru/slot
FRU hardware version	Hardware version of affected FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on affected FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment/ name
Attachment type	Specifically command output.	/mml/attachments/attachment/ type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/ mime
Command output text	Output of command automatically executed (see Table 54-3).	/mml/attachments/attachment/ atdata

Table 54-7 Inventory Event Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically "Inventory Update" Specific event names are listed in the "Event Triggers" section on page 54-21.	/mml/header/name
Message type	Specifically "Inventory Update".	/mml/header/type
Message group	Specifically "proactive".	/mml/header/group
Severity level	Severity level of inventory event is level 2 (seeTable 54-4).	/mml/header/level
Source ID	Product type for routing at Cisco. Specifically "MDS 9000"	/mml/header/source
Device ID	Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type*@*Sid*@*serial*, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /customerID
Contract ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: *type*@*Sid*@*serial*, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail

Table 54-7 Inventory Event Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the unit. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
FRU s/n	Serial number of FRU.	/mml/body/fru/serialNo
FRU part number	Part number of FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU.	/mml/body/fru/slot
FRU hardware version	Hardware version of FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment /name
Attachment type	Specifically command output.	/mml/attachments/attachment /type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment /mime
Command output text	Output of command automatically executed after event categories (see "Event Triggers" section on page 54-21).	/mml/attachments/attachment /atdata

Table 54-8 User-Generated Test Message Format

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically test message for test type message. Specific event names listed in the "Event Triggers" section on page 54-21).	/mml/header/name
Message type	Specifically "Test Call Home".	/mml/header/type
Message group	This field should be ignored by the receiving Call Home processing application, but may be populated with either "proactive" or "reactive".	/mml/header/group
Severity level	Severity level of message, test Call Home message (see Table 54-4).	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: *type@Sid@serial*, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C" identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /customerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: *type@Sid@serial*, where • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is "C" identifying the serial ID as a chassis serial number· • *serial* is the number identified by the Sid field. Example: "DS-C9509@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Switch that experienced the event.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact

Table 54-8 User-Generated Test Message Format (continued)

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Contact Email	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis. For example, 800-xxx-xxxx.	/mml/body/chassis/partNo
Command output text	Output of command automatically executed after event categories listed in Table 54-3.	/mml/attachments/attachment/atdata
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime
Attachment type	Specifically command output.	/mml/attachments/attachment/type
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name

CHAPTER 55

Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE, Fx, and TL ports) and their attached Nx ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

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

As of Cisco SAN-OS Release 3.1(2), FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, issued in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs. For devices that are zoned for IVR to be discovered with this command, they must have request domain_ID (RDI) enabled.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (Fx port). Hence your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In Figure 55-1 Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 55-1 FCSs in a VSAN Environment

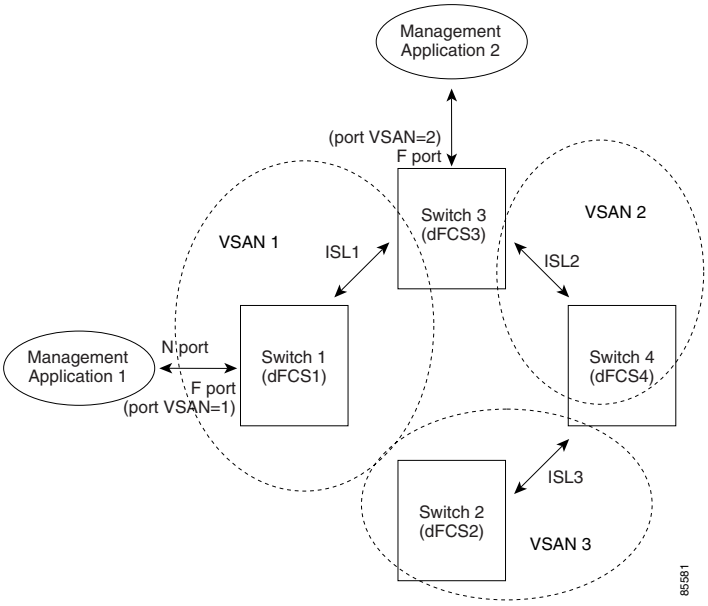

FCS Characteristics

FCSs have the following characteristics:

- FCSs support network management including the following:
  - N port management application can query and obtain information about fabric elements.
  - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs support TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

FCS Name Specification

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.

Note Set this command globally only if all switches in the fabric belong to the Cisco MDS 9000 Family.

To enable global checking of the platform name, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcs plat-check-global vsan 1`	Enables global checking of the platform name.
	`switch(config)# no fcs plat-check-global vsan 1`	Disables (default) global checking of the platform name.

To register platform attributes, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcs register` `switch(config-fcs-register)#`	Enters the FCS registration submode.
Step 3	`switch(config-fcs-register)# platform name SamplePlatform vsan 1` `switch(config-fcs-register-attrib)#`	Enters the FCS registration attributes submode.
	`switch(config-fcs-register)# no platform name SamplePlatform vsan 1` `switch(config-fcs-register)#`	Deletes a registered platform.
Step 4	`switch(config-fcs-register-attrib)# mgmt-addr 1.1.1.1`	Configures the platform management IPv4 address.
	`switch(config-fcs-register-attrib)# no mgmt-addr 1.1.1.1`	Deletes the platform management IPv4 address.
	`switch(config-fcs-register-attrib)# mgmt-addr 2001:0DB8:800:200C::417A`	Configures the platform management IPv6 address.
	`switch(config-fcs-register-attrib)# no mgmt-addr 2001:0DB8:800:200C::417A`	Deletes the platform management IPv6 address.
Step 5	`switch(config-fcs-register-attrib)# nwwn 11:22:33:44:55:66:77:88`	Configures the platform node name.
	`switch(config-fcs-register-attrib)# no nwwn 11:22:33:44:55:66:77:88`	Deletes the platform node name.
Step 6	`switch(config-fcs-register-attrib)# type 5`	Configures the fc-gs-3 defined platform type.
	`switch(config-fcs-register-attrib)# no type 5`	Deletes the configured type and reverts the switch to its factory default of unknown type.
Step 7	`switch(config-fcs-register-attrib)# exit`	Exits the FCS registration attributes submode.
Step 8	`switch(config-fcs-register)# exit` `switch(config)#`	Exits the FCS registration submode.

Displaying FCS Information

Use the **show fcs** commands to display the status of the WWN configuration (see Example 55-1 to 55-9).

Example 55-1 Displays FCS Local Database Information

```
switch# show fcs database
FCS Local Database in VSAN: 1
------------------------------
Switch WWN               : 20:01:00:05:30:00:16:df
Switch Domain Id         : 0x7f(127)
Switch Mgmt-Addresses    : snmp://172.22.92.58/eth-ip
                           http://172.22.92.58/eth-ip
Fabric-Name              : 20:01:00:05:30:00:16:df
Switch Logical-Name      : 172.22.92.58
Switch Information List  : [Cisco Systems*DS-C9509*0*20:00:00:05:30:00
Switch Ports:
-------------------------------------------------------------------
Interface  pWWN                      Type     Attached-pWWNs
-------------------------------------------------------------------
fc2/1      20:41:00:05:30:00:16:de  TE       20:01:00:05:30:00:20:de
fc2/2      20:42:00:05:30:00:16:de  Unknown  None
fc2/17     20:51:00:05:30:00:16:de  TE       20:0a:00:05:30:00:20:de

FCS Local Database in VSAN: 5
------------------------------
Switch WWN               : 20:05:00:05:30:00:12:5f
Switch Domain Id         : 0xef(239)
Switch Mgmt-Addresses    : http://172.22.90.171/eth-ip
                           snmp://172.22.90.171/eth-ip
                           http://10.10.15.10/vsan-ip
                           snmp://10.10.15.10/vsan-ip
Fabric-Name              : 20:05:00:05:30:00:12:5f
Switch Logical-Name      : 172.22.90.171
Switch Information List  : [Cisco Systems*DS-C9509**20:00:00:05:30:00:12:5e]
Switch Ports:
-------------------------------------------------------------------
Interface  pWWN                      Type     Attached-pWWNs
-------------------------------------------------------------------
fc3/1      20:81:00:05:30:00:12:5e  TE       22:01:00:05:30:00:12:9e
fc3/2      20:82:00:05:30:00:12:5e  TE       22:02:00:05:30:00:12:9e
fc3/3      20:83:00:05:30:00:12:5e  TE       22:03:00:05:30:00:12:9e
```

Example 55-2 Displays a List of All IEs for a Specific VSAN

```
switch# show fcs ie vsan 1
IE List for VSAN: 1
-------------------------------------------------------------------
IE-WWN                   IE-Type                          Mgmt-Id
-------------------------------------------------------------------
20:01:00:05:30:00:16:df  Switch (Local)                   0xfffc7f
20:01:00:05:30:00:20:df  Switch (Adjacent)                0xfffc64
[Total 2 IEs in Fabric]
```

Example 55-3 Displays Interconnect Element Object Information for a Specific nWWN

```
switch# show fcs ie nwwn 20:01:00:05:30:00:16:df vsan 1
IE Attributes
-------------
Domain-Id = 0x7f(127)
Management-Id = 0xfffc7f
Fabric-Name = 20:01:00:05:30:00:16:df
Logical-Name = 172.22.92.58
Management Address List =
        snmp://172.22.92.58/eth-ip
        http://172.22.92.58/eth-ip
Information List:
        Vendor-Name = Cisco Systems
        Model Name/Number = DS-C9509
        Release-Code = 0
```

Example 55-4 Displays Information for a Specific Platform

```
switch# show fcs platform name SamplePlatform vsan 1
Platform Attributes
-------------------
Platform Node Names:
        11:22:33:44:55:66:77:88
Platform Type = Gateway
Platform Management Addresses:
        1.1.1.1
```

Example 55-5 Displays a List of Platforms for a Specified VSAN

```
switch# show fcs platform vsan 1
Platform List for VSAN: 1
Platform-Names
--------------
SamplePlatform
[Total 1 Platforms in Fabric]
```

Example 55-6 Displays a List of Switch Ports in a Specified VSAN

```
switch# show fcs port vsan 24
Port List in VSAN: 24
                -- IE WWN: 20:18:00:05:30:00:16:df  --
------------------------------------------------------------------------
Port-WWN                Type      Module-Type            Tx-Type
------------------------------------------------------------------------
20:41:00:05:30:00:16:de  TE_Port   SFP with Serial Id      Shortwave Laser
20:51:00:05:30:00:16:de  TE_Port   SFP with Serial Id      Shortwave Laser
[Total 2 switch-ports in IE]
                -- IE WWN: 20:18:00:05:30:00:20:df  --
------------------------------------------------------------------------
Port-WWN                Type      Module-Type            Tx-Type
------------------------------------------------------------------------
20:01:00:05:30:00:20:de  TE_Port   SFP with Serial Id      Shortwave Laser
20:0a:00:05:30:00:20:de  TE_Port   SFP with Serial Id      Shortwave Laser
[Total 2 switch-ports in IE]
```

Example 55-7 Displays Port Information for a Specified pWWN

```
switch# show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24
Port Attributes
---------------
Port Type = TE_Port
Port Number = 0x1090000
Attached-Port-WWNs:
        20:0a:00:05:30:00:20:de
Port State = Online
```

Example 55-8 Displays FCS Statistics

```
switch# show fcs statistics
FCS Statistics for VSAN: 1
--------------------------
FCS Rx Get Reqs   :2
FCS Tx Get Reqs   :7
FCS Rx Reg Reqs   :0
FCS Tx Reg Reqs   :0
FCS Rx Dereg Reqs :0
FCS Tx Dereg Reqs :0
FCS Rx RSCNs      :0
...
FCS Statistics for VSAN: 30
--------------------------
FCS Rx Get Reqs   :2
FCS Tx Get Reqs   :2
FCS Rx Reg Reqs   :0
FCS Tx Reg Reqs   :0
FCS Rx Dereg Reqs :0
FCS Tx Dereg Reqs :0
FCS Rx RSCNs      :0
FCS Tx RSCNs      :0
...
```

Example 55-9 Displays Platform Settings for Each VSAN

```
switch# show fcs vsan
------------------------------
VSAN    Plat Check fabric-wide
------------------------------
0001    Yes
0010    No
0020    No
0021    No
0030    No
```

Default Settings

Table 55-1 lists the default FCS settings.

Table 55-1 Default FCS Settings

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

PART 9

Traffic Management

CHAPTER 56

Configuring Fabric Congestion Control and QoS

Fibre Channel Congestion Control (FCC) is a Cisco proprietary flow control mechanism that alleviates congestion on Fibre Channel networks.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

This chapter provides details on the QoS and FCC features provided in all switches. It includes the following sections:

FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. This section contains the following topics:

About FCC

The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see Figure 56-1).

Figure 56-1 FCC Mechanisms

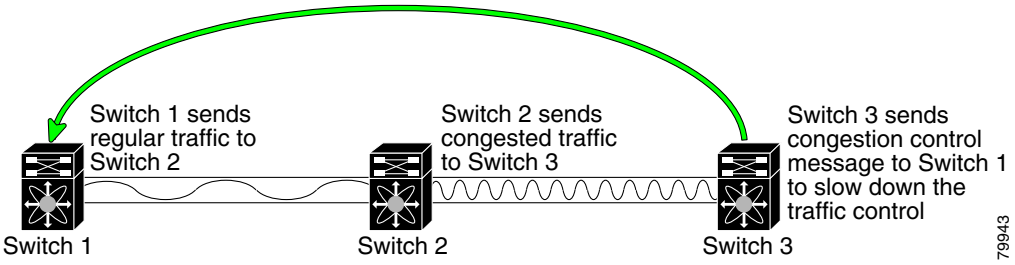

Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

Note FCC is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter.

FCC Process

When a node in the network detects congestion for an output port, it generates an edge quench message. These frames are identified by the Fibre Channel destination ID (DID) and the source ID. A switch from other vendors simply forwards these frames.

Any receiving switch in the Cisco MDS 9000 Family handles frames in one of these ways:

- It forwards the frame.
- It limits the rate of the frame flow in the congested port.

The behavior of the flow control mechanism differs based on the Fibre Channel DID:

- If the Fibre Channel DID is directly connected to one of the switch ports, the input rate limit is applied to that port.
- If the destination of the edge quench frame is a Cisco domain or the next hop is a Cisco MDS 9000 Family switch, the frame is forwarded.
- If neither of these mechanisms is true, then the frame is processed in the port going towards the FC DID.

All switches (including the edge switch) along the congested path process path quench frames. However, only the edge switch processes edge quench frames.

Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.

Tip If you enable FCC, be sure to enable it in all switches in the fabric.

To enable or disable the FCC feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcc`	Enables FCC in this switch.
	`switch(config)# no fcc`	Disables FCC in this switch (default).

Assigning FCC Priority

To assign FCC priority, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# fcc priority 2`	Defines the FCC priority threshold to have a priority of 2—0 is the lowest priority and 7 is the highest priority.

Displaying FCC Settings

Use the **show fcc** command to view FCC settings (see Example 56-1).

Example 56-1 Displays Configured FCC Information

```
switch# show fcc
fcc is disabled
fcc is applied to frames with priority up to 4
```

QoS

QoS implementation in the Cisco MDS 9000 Family follows the differentiated services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475.

All switches support the following types of traffic:

About Control Traffic

The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor's switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

Enabling or Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.

Tip We do not recommend disabling this feature as all critical control traffic is automatically assigned the lowest priority once you issue this command.

To disable the high priority assignment for control traffic, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# no qos control priority 0`	Enables the control traffic QoS feature.
	`switch(config)# qos control priority 0`	Disables the control traffic QoS feature.

Displaying Control Traffic Information

Use the **show qos statistics** command to view the current state of the QoS configuration for critical control traffic. This command displays the current QoS settings along with the number of frames marked high priority. The count is only for debugging purposes and cannot be configured (see Example 56-2).

Example 56-2 Displays Current QoS Settings

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted         = 8224
Current priority of FC control frames = 0    (0 = lowest; 7 = highest)
```

About Data Traffic

Online transaction processing (OLTP), which is a low volume, latency sensitive application, requires quick access to requested information. Backup processing application require high bandwidth but are not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically—they experience similar latency and are allocated similar bandwidths. The QoS feature in the Cisco MDS 9000 Family switches provides these guarantees.

Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing (see Figure 56-2).

Figure 56-2 Prioritizing Data Traffic

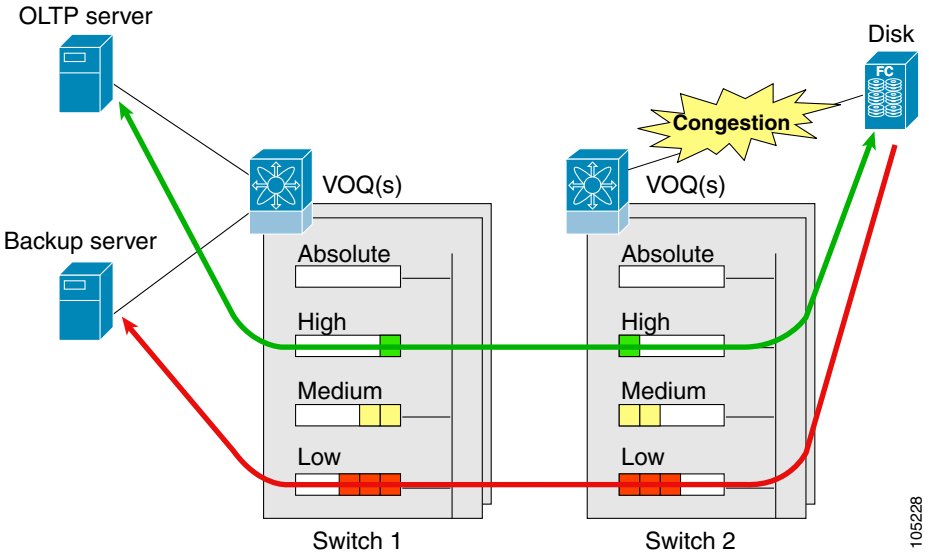

In Figure 56-2, the OLTP traffic arriving at Switch 1 is marked with a high priority level of throughput classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output queue (VOQ).

A deficit weighted round robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately since the high priority queue is not congested. The scheduler assigns its priority over the backup traffic in the low priority queue.

Note When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.

Tip To achieve this traffic differentiation, be sure to enable FCC (see the "Enabling FCC" section on page 56-2).

VSAN Versus Zone-Based QoS

While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch, both configurations have significant differences. Table 56-1 highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 56-1 QoS Configuration Differences

VSAN-Based QoS	Zone-Based QoS
If you configure the active zone set on a given VSAN and also configure QoS parameters in any of the member zones, you cannot associate the policy map with the VSAN.	You cannot activate a zone set on a VSAN that already has a policy map associated.
If the same flow is present in two class maps associated to a policy map, the QoS value of the class map attached first takes effect.	If the same flow is present in two zones in a given zone set with different QoS values, the higher QoS value is considered.
—	During a zone merge, if the Cisco SAN-OS software detects a mismatch for the QoS parameter, the link is isolated.
Takes effect only when QoS is enabled.	Takes effect only when QoS is enabled.

See the "About Zone-Based Traffic Priority" section on page 23-18 for details on configuring a zone-based QoS policy.

Configuring Data Traffic

To configure QoS, follow these steps:.

Step 1 Enable the QoS feature.

Step 2 Create and define class maps.

Step 3 Define service policies.

Step 4 Apply the configuration.

QoS Initiation for Data Traffic

By default, the QoS data traffic feature is disabled for data traffic. To configure QoS for data traffic, you must first enable the data traffic feature in the switch.

Tip QoS is supported in interoperability mode—its effectiveness depends on the location of Cisco MDS 9000 Family switches in the fabric relative to the location of the source or destination of the prioritized devices.

To enable the QoS data traffic feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# qos enable`	Enables QoS. You can now configure data traffic parameters.
	`switch(config)# no qos enable`	Removes the currently applied QoS configuration and disables QoS. You can no longer configure data traffic parameters.

About Class Map Creation

Use the class map feature to create and define a traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the match-all option. Flow-based traffic uses one of the following values:

- WWN—The source WWN or the destination WWN.
- Fibre Channel ID (FC ID) —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).

  Note An SID or DID of 0x000000 is not allowed.

- Source interface—The ingress interface.

Tip The order of entries to be matched within a class map is not significant.

Creating a Class Map

Use the **class-map** command to create and define a traffic class with match criteria to identify traffic belonging to that class. Define each match criterion with one match statement from the class map configuration (`switch(config-cmap))` mode.

- Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID).

- Use the **input-interface** option to specify the ingress interface.
- Use the **destination-device-alias** option to specify the distributed device alias.

To create a class map, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos class-map MyClass` `switch(config-cmap)#`	Creates a class map called MyClass and places you in the class-map submode to match all criteria specified for this class.
	`switch(config)# qos class-map MyClass match-all` `switch(config-cmap)#`	Specifies a logical AND operator for all matching statements in this class. If a frame matches all (default) configured criteria, it qualifies for this class. This is the default.
	`switch(config)# qos class-map MyClass match-any` `switch(config-cmap)#`	Specifies a logical OR operator for all matching statements in this class. If a frame matches any one configured criteria, it qualifies for this class.
Step 2	`switch(config-cmap)# match destination-address 0x12ee00`	Specifies a destination address match for frames with the specified destination FC ID.
	`switch(config-cmap)# match source-address 0x6d1090 mask 0xFFFFFF`	Specifies a source address and mask match for frames with the specified source FC ID.
Step 3	`switch(config-cmap)# match destination-wwn 20:01:00:05:30:00:28:df`	Specifies a destination WWN to match frames.
	`switch(config-cmap)# match source-wwn 23:15:00:05:30:00:2a:1f`	Specifies a source WWN to match frames.
Step 4	`switch(config-cmap)# match destination-device-alias DocDeviceAlias`	Specifies a destination device alias to match frames.
Step 5	`switch(config-cmap)# match input-interface fc 2/1`	Specifies a source interface to match frames.
Step 6	`switch(config-cmap)# no match input-interface fc 3/5`	Removes a match based on the specified source interface.

About Service Policy Definition

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. The policy map name is restricted to 63 alphanumeric characters.

As an alternative, you can map a class map to a differentiated services code point (DSCP).The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.

Note Refer to http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses for further information on implementing QoS DSCP values.

Note Class maps are processed in the order in which they are configured in each policy map.

Specifying Service Policies

To specify a service policy, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos policy-map MyPolicy` `switch(config-pmap)#`	Creates a policy map called MyPolicy and places you in the policy-map submode.
	`switch(config)# no qos policy-map OldPolicy` `switch(config)#`	Deletes the policy map called OldPolicy and places you in the policy-map submode.
Step 2	`switch(config-pmap)# class MyClass` `switch(config-pmap-c)#`	Specifies the name of a predefined class and places you at the policy-map submode for that class.
	`switch(config-pmap)# no class OldClass`	Removes the class map called OldClass from the policy map.
Step 3	`switch(config-pmap-c)# priority high`	Specifies the priority to be given for each frame matching this class.
	`switch(config-pmap-c)# no priority high`	Deletes a previously assigned priority and reverts to the default value of low.
Step 4	`switch(config-pmap-c)# dscp 2`	Specifies the DSCP value to mark each frame matching this class.
	`switch(config-pmap-c)# no dscp 60`	Deletes a previously assigned DSCP value and reverts to the factory default of 0.

About Service Policy Enforcement

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration is not enforced. You can only apply one policy map to a VSAN.

Note You can apply the same policy to a range of VSANs.

Applying Service Policies

To apply a service policy, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos service policy MyPolicy vsan 3`	Applies a configured policy to VSAN 3.
	`switch(config)# no qos service policy OldPolicy vsan 7`	Deletes a configured policy that was applied to VSAN 7.

About the DWRR Traffic Scheduler Queue

The Cisco SAN-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).

Table 56-2 describes the QoS behavior for Generation 1 and Generation 2 switching modules.

Table 56-2 QoS Behavior for Generation 1 and Generation 2 Switching Modules

Source Module Type	Destination Module Type	QoS Behavior Description
Generation 1	Generation 1	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other traffic share equal bandwidth.
Generation 1	Generation 2	QoS behavior reflects the DWRR configuration for traffic coming in through a given port and queued to the same egress port. All the other streams share equal bandwidth.
Generation 2	Generation 1	Bandwidth partitioning is equal for all the traffic.
Generation 2	Generation 2	QoS behavior reflects the DWRR weights configuration for all possible streams.

Changing the Weight in a DWRR Queue

To associate a weight with a DWRR queue, follow these steps:

	Command	Purpose
Step 1	`switch(config)# qos dwrr-q high weight 10`	Associates a relative weight (10) to a specified queue (default queue).
	`switch(config)# no qos dwrr-q low weight 51`	Restores the default weight of 20.

Displaying Data Traffic Information

The **show qos** commands display the current QoS settings for data traffic (see Examples 56-3 to 56-11).

Example 56-3 Displays the Contents of all Class Maps

```
switch# show qos class-map
qos class-map MyClass match-any
    match destination-wwn 20:01:00:05:30:00:28:df
    match source-wwn 23:15:00:05:30:00:2a:1f
    match input-interface fc2/1
qos class-map Class2 match-all
    match input-interface fc2/14
qos class-map Class3 match-all
    match source-wwn 20:01:00:05:30:00:2a:1f
```

Example 56-4 Displays the Contents of a Specified Class Map

```
switch# show qos class-map name MyClass
qos class-map MyClass match-any
    match destination-wwn 20:01:00:05:30:00:28:df
    match source-wwn 23:15:00:05:30:00:2a:1f
    match input-interface fc2/1
```

Example 56-5 Displays All Configured Policy Maps

```
switch# show qos policy-map
qos policy-map MyPolicy
    class MyClass
    priority medium
qos policy-map Policy1
    class Class2
    priority low
```

Example 56-6 Displays a Specified Policy Map

```
switch# show qos policy-map name MyPolicy
qos policy-map MyPolicy
    class MyClass
        priority medium
```

Example 56-7 Displays Scheduled DWRR Configurations

```
switch# show qos dwrr
qos dwrr-q high weight 50
qos dwrr-q medium weight 30
qos dwrr-q low weight 20
```

Example 56-8 Displays All Applied Policy Maps

```
switch# show qos service policy
qos service policy MyPolicy vsan 1
qos service policy Policy1 vsan 4
```

Example 56-9 Displays the Policy Map Associated with a Specified VSAN

```
switch# show qos service policy vsan 1
qos policy-map pmap1
   class cmap1
       priority medium
   class cmap2
       priority high
```

Example 56-10 Displays the Class Map Associated with a Specified Interface

```
switch# show qos service policy interface fc3/10
qos policy-map pmap1
   class cmap3
       priority high
   class cmap4
       priority low
```

Example 56-11 Displays QoS Statistics

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted         = 137679
Current priority of FC control frames = 7    (0 = lowest; 7 = highest)
```

Example Configuration

This section describes a configuration example for the application illustrated in Figure 56-3.

Figure 56-3 Example Application for Traffic Prioritization

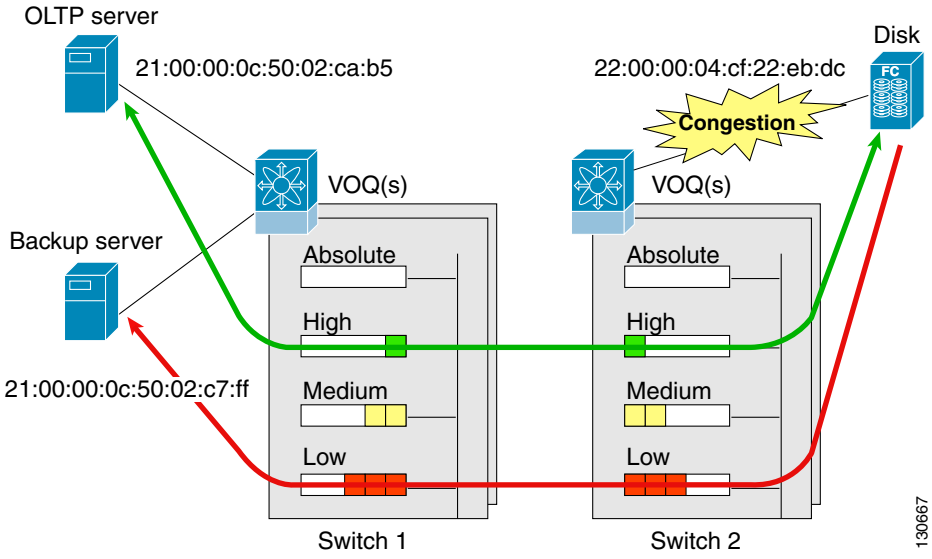

Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires faster response because transaction processing is a low latency application.

The point of congestion is the link between Switch 2 and the disk, for traffic from the switch to the disk. The return path is largely uncongested as there is little backup traffic on this path.

Service differentiation is needed at Switch 2 to prioritize the OLTP-server-to-disk traffic higher than the backup-server-to-disk traffic.

To configure traffic prioritization for the example application, follow these steps:

Step 1 Create the class maps.

```
Switch 2# config t
Switch 2(config)# qos class-map jc1 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc2 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Create the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc1
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc2
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Assign the service policy.

```
Switch 2(config)# qos service policy jp1 vsan 1
```

Step 4 Assign the weights for the DWRR queues.

```
Switch 2(config)# qos dwrr-q high weight 50
Switch 2(config)# qos dwrr-q medium weight 30
Switch 2(config)# qos dwrr-q low weight 20
```

Step 5 Repeat Step 1 through Step 4 on Switch 1 to address forward path congestion at both switches.

Congestion could occur anywhere in the example configuration. To address congestion of the return path at both switches, you need to create two more class maps and include them in the policy map as follows:

Step 1 Create two more class maps.

```
Switch 2(config)# qos class-map jc3 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc4 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Assign the class maps to the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc3
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc4
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Repeat Step 1 through Step 2 on Switch 1 to address return path congestion at both switches.

Ingress Port Rate Limiting

A port rate limiting feature helps control the bandwidth for individual Fibre Channel ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a Fibre Channel port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.

Note Port rate limiting can only be configured on Cisco MDS 9100 Series switches, Cisco MDS 9216i switches, and MPS-14/2 modules.

This feature can only be configured if the QoS feature is enabled and if this configuration is performed on a Cisco MDS 9100 series switch, Cisco MDS 9216i switch, or MPS-14/2 module.

To configure the port rate limiting value, follow these steps.

	Command	Purpose
Step 1	`switch # config t` `switch(config)#`	Enters the configuration mode.
Step 2	`switch(config)# interface fc 1/1`	Selects the interface to specify the ingress port rate limit.
Step 3	`switch(config-if)# switchport ingress-rate 50`	Configures a 50% port rate limit for the selected interface.
	`switch(config-if)# no switchport ingress-rate 50`	Reverts a previously configured rate to the factory default of 100%.

Default Settings

Table 56-3 lists the default settings for FCC, QoS, and rate limiting features.

Table 56-3 Default FCC, QoS, and Rate Limiting Settings

Parameters	Default
FCC protocol	Disabled.
QoS control traffic	Enabled.
QoS data traffic	Disabled.
Zone-based QoS priority	Low.
Rate limit	100%.

CHAPTER 57

Configuring Port Tracking

The port tracking feature is unique to the Cisco MDS 9000 Family of switches. 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.

This chapter includes the following sections:

About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information (see the "Common Information Model" section on page 29-1 and "About RSCN Information" section on page 26-8).

In Figure 57-1, 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 57-1 Traffic Recovery Using Port Tracking

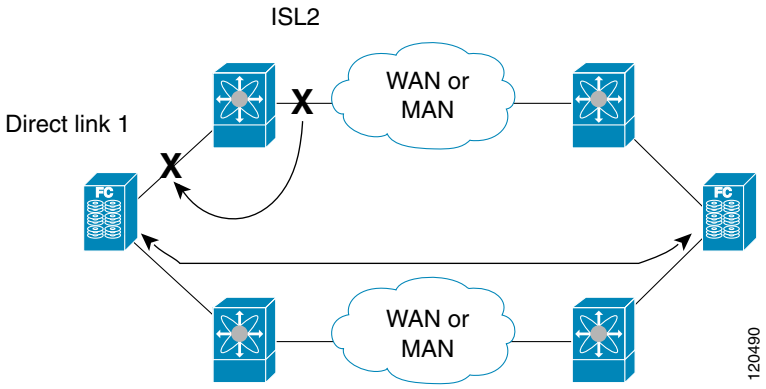

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco SAN-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter.

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports.
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS 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 fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.

This section includes the following topics:

About Port Tracking

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 tracked 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 the port up when required.

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-track enable`	Enables port tracking.
	`switch(config)# no port-track enable`	Removes the currently applied port tracking configuration and disables port tracking.

About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6` `switch(config-if)#`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports. **Note** This link symbolizes the direct link (1) in Figure 57-1.
Step 3	`switch(config-if)# port-track interface port-channel 1`	Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (2) in Figure 57-1.
	`switch(config-if)# no port-track interface port-channel 1`	Removes the port tracking configuration that is currently applied to interface fc8/6.

About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In Figure 57-2, 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 57-2 Traffic Recovery Using Port Tracking

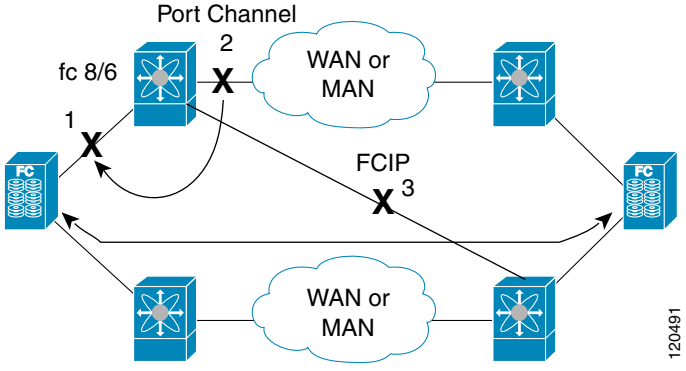

Tracking Multiple Ports

To track multiple ports, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports. **Note** This link symbolizes the direct link (1) in Figure 57-2.
Step 3	`switch(config-if)# port-track interface port-channel 1`	Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (2) in Figure 57-2.
Step 4	`switch(config-if)# port-track interface fcip 5`	Tracks interface fc8/6 with interface fcip 5. When FCIP 5 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (3) in Figure 57-2.

About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

Tip	The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.
Step 3	`switch(config-if)# port-track interface port-channel 1 vsan 2`	Enables tracking of the PortChannel in VSAN 2.
	`switch(config-if)# no port-track interface port-channel 1 vsan 2`	Removes the VSAN association for the linked port. The PortChannel link remains in effect.

AboutForceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

To forcefully shut down a tracked port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/5`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.
Step 3	`switch(config-if)# port-track force-shut`	Forcefully shuts down the tracked port.
	`switch(config-if)# no port-track force-shut`	Removes the port shutdown configuration for the tracked port.

Displaying Port Tracking Information

The **show** commands display the current port tracking settings for the Cisco MDS switch (see Examples 57-1 to 57-4).

Example 57-1 Displays the Linked and Tracked Port Configuration

```
switch# show interface
...
fc8/6 is down (All tracked ports down)     <-----------------------------------Linked port
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 21:c6:00:05:30:00:37:1e
    Admin port mode is auto, trunk mode is on
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface port-channel 1 vsan 2 (trunking) <-----Tracked port
    Port tracked with interface fcip 5 <--------------------------------------Tracked port
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      269946 frames input, 22335204 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      205007 frames output, 10250904 bytes
```

```
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      2 output OLS, 2 LRR, 0 NOS, 1 loop inits
      0 receive B2B credit remaining
      0 transmit B2B credit remaining
...
```

Example 57-2 Displays a Tracked Port Configuration for a Fibre Channel Interface

```
switch# show interface fc1/1
fc1/1 is down (Administratively down)
    Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
    Port WWN is 20:01:00:05:30:00:0d:de
    Admin port mode is FX
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface fc1/2 (down)
    Port tracked with interface port-channel 1 vsan 2 (down)
    Port tracked with interface fcip1 (down)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1 frames input, 128 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      1 frames output, 128 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
      0 receive B2B credit remaining
      0 transmit B2B credit remaining
```

Example 57-3 Displays a Tracked Port Configuration for a PortChannel Interface

```
switch# show interface port-channel 1
port-channel 1 is down (No operational members)
    Hardware is Fibre Channel
    Port WWN is 24:01:00:05:30:00:0d:de
    Admin port mode is auto, trunk mode is on
    Port vsan is 2
    Linked to 1 port(s)
      Port linked to interface fc1/1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    No members
```

Example 57-4 Displays a Forced Shutdown Configuration

```
switch# show interface fc 1/5
```

```
fc1/5 is up
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:05:00:05:30:00:47:9e
    Admin port mode is F
    Port mode is F, FCID is 0x710005
    Port vsan is 1
    Speed is 1 Gbps
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Port track mode is force_shut <--this port remains shut even if the tracked port is back up
```

Default Port Tracking Settings

Table 57-1 lists the default settings for port tracking parameters.

Table 57-1 Default Port Tracking Parameters

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

PART 10

Troubleshooting

CHAPTER 58

Troubleshooting Your Fabric

This chapter describes basic troubleshooting methods used to resolve issues with switches. This chapter includes the following sections:

fctrace

The fctrace feature allows you to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

You can invoke fctrace by providing the FC ID, the N port, or the NL port WWN, or the device alias of the destination. The frames are routed normally as long as they are forwarded through TE ports.

Once the frame reaches the edge of the fabric (the F port or FL port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of failure.

Note The fctrace feature works only on TE ports. Make sure that only TE ports exist in the path to the destination. In case there is an E port in the path, the fctrace frame is dropped by that switch. Also, fctrace times out in the originator, and path discovery does not start.

Tip You cannot use the fctrace feature in a locally configured VSAN interface (IPFC interface), but you can trace the route to a VSAN interface configured in other switches.

To perform a fctrace operation, follow this step:

	Command	Purpose

Step 1	`switch# fctrace fcid 0xd70000 vsan 1`
	`Route present for : 0xd70000`
	`20:00:00:0b:46:00:02:82(0xfffcd5)`
	`Timestamp Invalid.`
	`20:00:00:05:30:00:18:db(0xfffcd7)`
	`Timestamp Invalid.`
	`20:00:00:05:30:00:18:db(0xfffcd7)`
		Invokes fctrace for the specified FC ID of the destination N port.

	`switch# fctrace pwwn 21:00:00:e0:8b:06:d9:1d vsan 1 timeout 5`
	`Route present for : 21:00:00:e0:8b:06:d9:1d`
	`20:00:00:0b:46:00:02:82(0xfffcd5)`
	`Timestamp Invalid.`
	`20:00:00:05:30:00:18:db(0xfffcd7)`
	`Timestamp Invalid.`
	`20:00:00:05:30:00:18:db(0xfffcd7)`
		Invokes fctrace using the pWWN of the destination N port.
		By default the period to wait before timing out is 5 seconds, The range is from one through 10 seconds.

	`switch# fctrace device-alias disk1 v 1`
	`Route present for : 22:00:00:0c:50:02:ce:f8`
	`20:00:00:05:30:00:31:1e(0xfffca9)`
		Invokes fctrace using the device alias of the destination N port.

fcping

The fcping feature verifies reachability of a node by checking its end-to-end connectivity. You can invoke the fcping feature by providing the FC ID, the destination port WWN, or the device alias information.

To perform a fcping operation, follow these steps:

	Command	Purpose

Step 1	`switch# fcping fcid 0xd70000 vsan 1`
	`28 bytes from 0xd70000 time = 730 usec`
	`28 bytes from 0xd70000 time = 165 usec`
	`28 bytes from 0xd70000 time = 262 usec`
	`28 bytes from 0xd70000 time = 219 usec`
	`28 bytes from 0xd70000 time = 228 usec`
	
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 165/270/730 usec`
		Invokes fcping for the specified pWWN or the FC ID of the destination. By default, five frames are sent.

	`switch# fcping fcid 0xd70000 vsan 1 count 10`
	`28 bytes from 0xd70000 time = 730 usec`
	`28 bytes from 0xd70000 time = 165 usec`
	`28 bytes from 0xd70000 time = 262 usec`
	`28 bytes from 0xd70000 time = 219 usec`
	`28 bytes from 0xd70000 time = 228 usec`
	`28 bytes from 0xd70000 time = 230 usec`
	`28 bytes from 0xd70000 time = 230 usec`
	`28 bytes from 0xd70000 time = 225 usec`
	`28 bytes from 0xd70000 time = 229 usec`
	`28 bytes from 0xd70000 time = 183 usec`
	
	`10 frames sent, 10 frames received, 0 timeouts`
	`Round-trip min/avg/max = 165/270/730 usec`
		Sets the number of frames to be sent using the count option. The range is from 0 through 2147483647. A value of 0 pings forever.

	`switch# fcping fcid 0xd500b4 vsan 1 timeout 10`
	`28 bytes from 0xd500b4 time = 1345 usec`
	`...`
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 340/581/1345 usec`
		Sets the timeout value. The default period to wait is 5 seconds. The range is from 1 through 10 seconds.

	`switch# fcping device-alias disk1 vsan 1`
	`28 bytes from 22:00:00:0c:50:02:ce:f8 time = 1883 usec`
	`28 bytes from 22:00:00:0c:50:02:ce:f8 time = 493 usec`
	`28 bytes from 22:00:00:0c:50:02:ce:f8 time = 277 usec`
	`28 bytes from 22:00:00:0c:50:02:ce:f8 time = 391 usec`
	`28 bytes from 22:00:00:0c:50:02:ce:f8 time = 319 usec`
	
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 277/672/1883 usec`
		Invokes fcping for the specified device alias of the destination.

Step 2	`switch# fcping fcid 0x010203 vsan 1`
	`No response from the N port.`
	
	`switch# fcping pwwn 21:00:00:20:37:6f:db:dd vsan 1`
	`28 bytes from 21:00:00:20:37:6f:db:dd time = 1454 usec`
	`...`
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 364/784/1454 usec`
		Issues a `No response from the N port` message even when the N port or NL port is active. This is due to resource exhaustion at the N port or NL port.
		
		Retry the command a few seconds later.

Verifying Switch Connectivity

You can verify connectivity to a destination switch.

Note The FC ID variable used in this procedure is the domain controller address; it is not a duplication of the domain ID.

To verify connectivity to a destination switch, follow these steps:

	Command	Purpose

Step 1	`switch# show fcdomain domain-list vsan 200`
	`Number of domains: 7`
	`Domain ID WWN`
	`--------- -----------------------`
	`0x01(1) 20:c8:00:05:30:00:59:df [Principal]`
	`0x02(2) 20:c8:00:0b:5f:d5:9f:c1`
	`0x6f(111) 20:c8:00:05:30:00:60:df`
	`0xda(218) 20:c8:00:05:30:00:87:9f [Local]`
	`0x06(6) 20:c8:00:0b:46:79:f2:41`
	`0x04(4) 20:c8:00:05:30:00:86:5f`
	`0x6a(106) 20:c8:00:05:30:00:f8:e3`
		Displays the destination switch's domain ID.
		To obtain the domain controller address, concatenate the domain ID with FFFC. For example, if the domain ID is 0xda(218), the concatenated ID is 0xfffcda.

Step 2	`switch# fcping fcid 0xFFFCDA vsan 200`
	`28 bytes from 0xFFFCDA time = 298 usec`
	`28 bytes from 0xFFFCDA time = 260 usec`
	`28 bytes from 0xFFFCDA time = 298 usec`
	`28 bytes from 0xFFFCDA time = 294 usec`
	`28 bytes from 0xFFFCDA time = 292 usec`
	
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 260/288/298 usec`
		Verifies reachability of the destination switch by checking its end-to-end connectivity.

Cisco Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. Existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, to snoop traffic, the existing analyzers disrupt the traffic on the link while the analyzer is inserted into the link.

The Cisco MDS 9000 Family switches support protocol analysis within a storage network with the Cisco Fabric Analyzer. You can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

The Cisco Fabric Analyzer is based on two popular public-domain software applications:

- libpcap—See http://www.tcpdump.org.
- Ethereal—See http://www.ethereal.com.

Note The Cisco Fabric Analyzer is useful in capturing and decoding control traffic, not data traffic. It is suitable for control path captures, and is not intended for high-speed data path captures.

This section includes the following topics:

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer consists of two separate components (see Figure 58-1):

- Software that runs on the Cisco MDS 9000 Family switch and supports two modes of capture:
  - A text-based analyzer that supports local capture and decodes captured frames
  - A daemon that supports remote capture
- GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco MDS 9000 Family switch.

Figure 58-1 Cisco Fabric Analyzer Usage

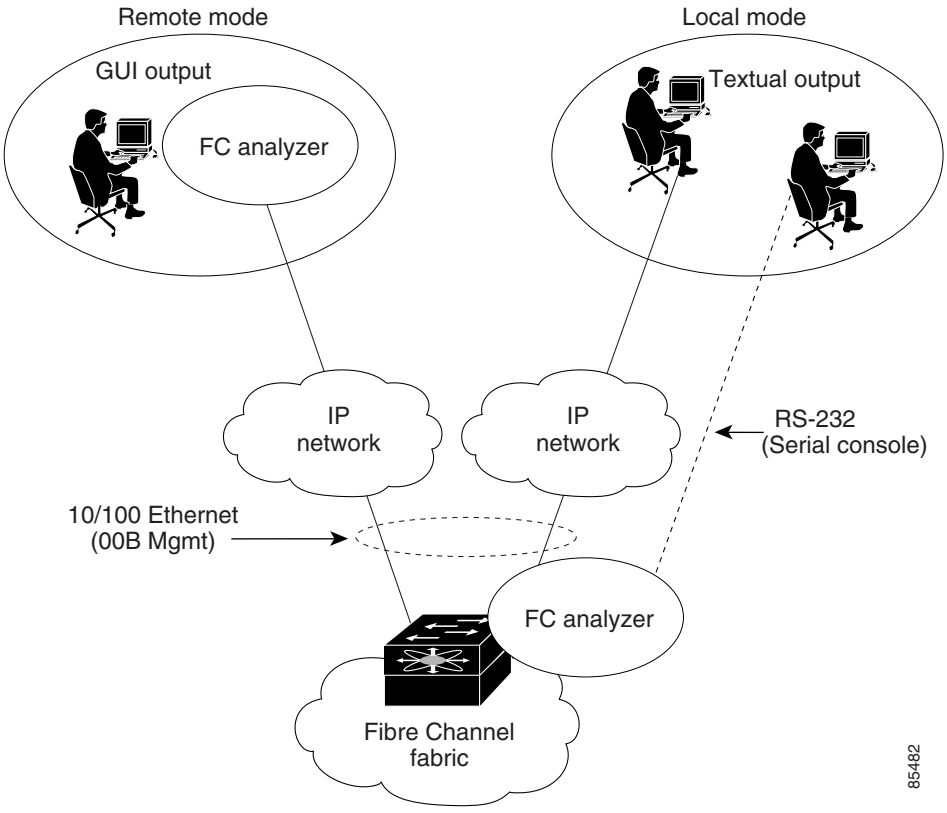

Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 Family switch. It is a fully functional decoder that is useful for quick debug purposes or for use when the remote capture daemon is not enabled. Additionally, because this tool is accessed from within the Cisco MDS 9000 Family switch, it is protected by the roles-based policy that limits access in each switch.

See the "Capturing Frames Locally" section on page 58-7.

Remote Capture Daemon

This daemon is the server end of the remote capture component. The Ethereal analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two endpoints, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents an unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on firewall restrictions.

- Passive mode (default)—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- Active mode—The switch initiates the connection to a configured host—one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end—on Ethereal, not on the switch.

See the "Sending Captures to Remote IP Addresses" section on page 58-8.

GUI-Based Client

The Ethereal software runs on a host, such as a PC or workstation, and communicates with the remote capture daemon. This software is available in the public domain from http://www.ethereal.com. The Ethereal GUI front-end supports a rich interface such as a colorized display, graphical assists in defining filters, and specific frame searches. These features are documented on Ethereal's web site.

While remote capture through Ethereal supports capturing and decoding Fibre Channel frames from a Cisco MDS 9000 Family switch, the host running Ethereal does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

See the "Displaying Captured Frames" section on page 58-10.

Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer to perform one of two captures.

- Local capture—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby. Launches the textual version on the fabric analyzer directly on the console screen. The capture can also be saved on the local file system.
- Remote capture—The command setting to enable a remote capture can be saved to persistent storage. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

Capturing Frames Locally

To capture frames locally, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
	Note The options within Step 2 may be performed in any order.	

Step 2	`switch(config)# fcanalyzer local`
	`Capturing on eth2`
	`switch(config)#`
		Begins capturing the frames locally (supervisor module).

	`switch(config)# fcanalyzer local brief`
	`Capturing on eth2`
	`switch(config)#`
		Displays the protocol summary in a brief format.

	`switch(config)# fcanalyzer local display-filter SampleF` `Capturing on eth2`	Displays the filtered frames.

	`switch(config)# fcanalyzer local limit-frame-size 64`
	`Capturing on eth2`
	`switch(config)#`
		Limits the size of the frame capture to the first 64 bytes. The allowed range is 64 to 65536 bytes.

	`switch(config)# fcanalyzer local limit-captured-frames 10`
	`Capturing on eth2`
	`switch(config)#`
		Limits the number of frames captured to 10. The allowed range is 0 to 2147483647 frames and the default is 100 frames. Use 0 if you do not want to limit the number of captured frames.

	Note Press **Ctrl-c** to stop a capture. Otherwise, the capture stops automatically after capturing 100 frames. You can change this default using the **fcanalyzer local limit-captured-frames** ***number*** command.	

	Command	Purpose

Step 3	`switch(config)# fcanalyzer local write volatile:sample`
	`Capturing on eth2`
	`switch(config)#`
		Saves the captured frames to a specified file (sample) in the volatile: directory.
		Note Optionally, you can save the specified file to the slot0: directory.

Note The final file name is the capture file called either SampleFile_00000_*dateandtime* or SampleFile_00001_*dateandtime*.
For example, “SampleFile_00000_20021110223833” or “SampleFile_00001_20021110243833”.
The maximum size of a file that can be written to is 10 MB.

Sending Captures to Remote IP Addresses

Caution You must use the eth2 interface to capture control traffic on a supervisor module.

To capture frames remotely using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcanalyzer remote 10.21.0.3`	Configures the remote IPv4 address (10.21.0.3) to which the captured frames are sent.
	`switch(config)# fcanalyzer remote 10.21.0.3 active`	Enables active mode (passive is the default) with the remote host. Ethereal is assumed to be running when the capture is performed. The switch tries to connect forever unless a capture stop instruction is sent from Ethereal.
	`switch(config)# fcanalyzer remote 10.21.0.3 active 1`	Enables the active mode for a specified port. The valid port range is 1 to 65535.

To capture frames remotely using IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# fcanalyzer remote 2001:0DB8:800:200C::417A`	Configures the remote IPv6 address to which the captured frames are sent.
	`switch(config)# fcanalyzer remote 2001:0DB8:800:200C::417A active`	Enables active mode (passive is the default) with the remote host. Ethereal is assumed to be running when the capture is performed. The switch tries to connect forever unless a capture stop instruction is sent from Ethereal.
	`switch(config)# fcanalyzer remote 2001:0DB8:800:200C::417A active 1`	Enables the active mode for a specified port. The valid port range is 1 to 65535.

To capture remote traffic, use one of the following options:

- The capture interface can be specified in Ethereal as the remote device:

  ```
  rpcap://<ipaddress or switch hostname>/eth2
  ```

  For example:

  ```
  rpcap://cp-16/eth2
  rpcap://17.2.1.1/eth2
  ```

- The capture interface can be specified either in the capture dialog box or by using the -i option at the command line when invoking Ethereal.

  ```
  ethereal -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
  ```

  For example:

  ```
  ethereal -i rpcap://172.22.1.1/eth2
  ```

  or

  ```
  ethereal -i rpcap://customer-switch.customer.com/eth2
  ```

Note For example, in a Windows 2000 setup, click **Start** on your desktop and select **Run**. In the resulting Run window, type the required command line option in the Open field.

Clearing Configured fcanalyzer Information

Use the **clear fcanalyzer** command to clear the entire list of configured hosts. Note that the existing connections are not terminated.

Displaying Configured Hosts

Use the **show fcanalyzer** command to display the list of hosts configured for a remote capture. See Example 58-1.

Example 58-1 Displays Configured Hosts

```
switch# show fcanalyzer
PassiveClient = 10.21.0.3
PassiveClient = 10.21.0.3
ActiveClient = 10.21.0.3, DEFAULT
```

Note The DEFAULT in the ActiveClient line indicates that the default port is used.

Displaying Captured Frames

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view Exchange Link Protocol (ELP) request frames. This feature only limits the captured view—it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already documented in the Ethereal web site (http://www.ethereal.com). Some examples of how you can use this feature are as follows:

- To view all packets in a specified VSAN, use this expression:

  ```
  mdshdr.vsan == 2
  ```

- To view all SW_ILS frames, use this expression:

  ```
  fcswils
  ```

- To view class F frames, use this expression:

  ```
  mdshdr.sof == SOFf
  ```

- To view all FSPF frames, use this expression:

  ```
  swils.opcode == HLO || swils.opcode == LSU || swils.opcode == LSA
  ```

- To view all FLOGI frames, use this expression:

  ```
  fcels.opcode == FLOGI
  ```

- To view all FLOGI frames in VSAN 1, use this expression:

  ```
  fcels.opcode == FLOGI && mdshdr.vsan == 2
  ```

- To view all name server frames, use this expression:

  ```
  dNS
  ```

Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.

Note This GUI-assisted feature is part of Ethereal and you can obtain more information from http://www.ethereal.com.

Examples of Display Filters

Some examples of using display filters with the Fabric Analyzer local are provided in this section. The **brief** option is used in all examples to restrict the size of the output. See Example 58-2.

Example 58-2 Displays Only Fabric Login Server Traffic on VSAN 1

```
switch(config)# fcanalyzer local brief display-filter
(mdshdr.vsan==0x01)&&((fc.d_id==\"ff.ff.fe\"\|\|fc.s_id==\"ff.ff.fe\"))
Capturing on eth2
8.904145   00.00.00 -> ff.ff.fe   FC ELS 1    0x28f8 0xffff  0x3 ->  0xf FLOGI
8.918164   ff.ff.fe -> 79.03.00   FC ELS 1    0x28f8 0x12c6 0xff ->  0x0 ACC (FLOGI)
```

You can trace all frames to and from a particular N port device. For example, you can observe RSCNs from the Fabric Controller and registration, and/or you can query requests to the name server. See Example 58-3.

Note The filter requires prior knowledge of the FC ID that is assigned to the N port. Issue the **show flogi database interfac**e command before running fcanalyzer to obtain the FC ID. In this example, the N port FC ID is 79.03.00.

Example 58-3 Displays All Traffic for a Particular N Port on VSAN 1

```
switch(config)# fcanalyzer local brief
display-filter(mdshdr.vsan==0x01)&&((fc.d_id==\"79.03.00\"\|\|fc.s_id==\"79.03.00\"))
Capturing on eth2
8.699162   ff.ff.fe -> 79.03.00   FC ELS 1    0x35b8 0x148e 0xff ->  0x0 ACC (FLOGI)
8.699397   79.03.00 -> ff.ff.fc   FC ELS 1    0x35d0 0xffff  0x3 ->  0xf PLOGI
8.699538   ff.ff.fc -> 79.03.00   FC ELS 1    0x35d0 0x148f 0xff ->  0x0 ACC (PLOGI)
8.699406   79.03.00 -> ff.ff.fd   FC ELS 1    0x35e8 0xffff  0x3 ->  0xf SCR
8.700179   79.03.00 -> ff.ff.fc   dNS    1    0x3600 0xffff  0x3 ->  0xf GNN_FT
8.702446   ff.ff.fd -> 79.03.00   FC ELS 1    0x35e8 0x1490 0xff ->  0x0 ACC (SCR)
8.704210   ff.ff.fc -> 79.03.00   dNS    1    0x3600 0x1491 0xff ->  0x0 ACC (GNN_FT)
8.704383   79.03.00 -> ff.ff.fc   dNS    1    0x3618 0xffff  0x3 ->  0xf GPN_ID
8.707857   ff.ff.fc -> 79.03.00   dNS    1    0x3618 0x1496 0xff ->  0x0 ACC (GPN_ID)
```

The VSAN ID is specified in hex. See Example 58-4.

Example 58-4 Displays All Traffic for a Specified VSAN

```
switch(config)# fcanalyzer local brief display-filter mdshdr.vsan==0x03e7
Capturing on eth2
12.762577   ff.ff.fd -> ff.ff.fd   SW_ILS 999   0xb2c 0xffff  0x1 ->  0xf HLO
12.762639   ff.ff.fd -> ff.ff.fd   FC     999   0xb2c  0xd32 0xff ->  0x0 Link Ctl, ACK1
13.509979   ff.ff.fd -> ff.ff.fd   SW_ILS 999   0xd33 0xffff 0xff ->  0x0 HLO
13.510918   ff.ff.fd -> ff.ff.fd   FC     999   0xd33  0xb2d  0x1 ->  0xf Link Ctl, ACK1
14.502391   ff.fc.64 -> ff.fc.70   SW_ILS 999   0xd34 0xffff 0xff ->  0x0 SW_RSCN
14.502545   ff.ff.fd -> 64.01.01   FC ELS 999   0xd35 0xffff 0xff ->  0x0 RSCN
14.502804   64.01.01 -> ff.ff.fd   FC ELS 999   0xd35  0x215  0x0 ->  0xf ACC (RSCN)
14.503387   ff.fc.70 -> ff.fc.64   FC     999   0xd34  0xb2e  0x1 ->  0xf Link Ctl, ACK1
14.503976   ff.fc.70 -> ff.fc.64   SW_ILS 999   0xd34  0xb2e  0x1 ->  0xf SW_ACC (SW_RSCN)
14.504025   ff.fc.64 -> ff.fc.70   FC     999   0xd34  0xb2e 0xff ->  0x0 Link Ctl, ACK1
```

By excluding FSPF `hellos` and `ACK1`, you can focus on the frames of interest. See Example 58-5.

Example 58-5 Displays All VSAN 1 Traffic Excluding FSPF Hellos and ACK1 Frames.

```
switch(config)# fcan lo bri dis
(mdshdr.vsan==0x01)&&not((swils.opcode==0x14)or(fc.r_ctl==0xc0))
Capturing on eth2
10.589934   ff.fc.79 -> ff.fc.7a   FC-FCS 1    0x1b23 0xffff 0xff ->  0x0 GCAP
10.591253   ff.fc.7a -> ff.fc.79   FC-FCS 1    0x1b23 0x2f70  0x4 ->  0xf MSG_RJT (GCAP)
25.277981   ff.fc.79 -> ff.fc.7a   SW_ILS 1    0x1b27 0xffff 0xff ->  0x0 SW_RSCN
25.278050   ff.fc.79 -> ff.fc.89   SW_ILS 1    0x1b28 0xffff 0xff ->  0x0 SW_RSCN
25.279232   ff.fc.89 -> ff.fc.79   SW_ILS 1    0x1b28 0xadd7  0x5 ->  0xf SW_ACC (SW_RSCN)
25.280023   ff.fc.7a -> ff.fc.79   Unzoned NS 1    0x3b2b 0xffff  0x5 ->  0xf GE_PT
25.280029   ff.fc.7a -> ff.fc.79   SW_ILS 1    0x1b27 0x2f71  0x4 ->  0xf SW_ACC (SW_RSCN)
25.282439   ff.fc.79 -> ff.fc.7a   dNS    1    0x3b2b 0x1b29 0xff ->  0x0 RJT (GE_PT)
38.249966   00.00.00 -> ff.ff.fe   FC ELS 1    0x36f0 0xffff  0x3 ->  0xf FLOGI
38.262622   ff.ff.fe -> 79.03.00   FC ELS 1    0x36f0 0x1b2b 0xff ->  0x0 ACC (FLOGI)
38.262844   79.03.00 -> ff.ff.fc   FC ELS 1    0x3708 0xffff  0x3 ->  0xf PLOGI
38.262984   ff.ff.fc -> 79.03.00   FC ELS 1    0x3708 0x1b2c 0xff ->  0x0 ACC (PLOGI)
38.262851   79.03.00 -> ff.ff.fd   FC ELS 1    0x3720 0xffff  0x3 ->  0xf SCR
38.263514   ff.fc.79 -> ff.fc.7a   SW_ILS 1    0x1b2e 0xffff 0xff ->  0x0 SW_RSCN
38.263570   ff.fc.79 -> ff.fc.89   SW_ILS 1    0x1b2f 0xffff 0xff ->  0x0 SW_RSCN
38.263630   79.03.00 -> ff.ff.fc   dNS    1    0x3738 0xffff  0x3 ->  0xf GNN_FT
38.263884   ff.ff.fd -> 79.03.00   FC ELS 1    0x3720 0x1b2d 0xff ->  0x0 ACC (SCR)
38.264066   ff.fc.89 -> ff.fc.79   SW_ILS 1    0x1b2f 0xaddf  0x5 ->  0xf SW_ACC (SW_RSCN)
38.264417   ff.fc.89 -> ff.fc.79   dNS    1    0xade0 0xffff  0x5 ->  0xf GE_ID
38.264585   ff.fc.79 -> ff.fc.89   dNS    1    0xade0 0x1b31 0xff ->  0x0 ACC (GE_ID)
38.265132   ff.ff.fc -> 79.03.00   dNS    1    0x3738 0x1b30 0xff ->  0x0 ACC (GNN_FT)
38.265210   ff.fc.7a -> ff.fc.79   Unzoned NS 1    0x3b2f 0xffff  0x5 ->  0xf GE_PT
38.265414   79.03.00 -> ff.ff.fc   dNS    1    0x3750 0xffff  0x3 ->  0xf GPN_ID
38.265502   ff.fc.7a -> ff.fc.79   SW_ILS 1    0x1b2e 0x2f73  0x4 ->  0xf SW_ACC (SW_RSCN)
38.267196   ff.fc.79 -> ff.fc.7a   dNS    1    0x3b2f 0x1b32 0xff ->  0x0 ACC (GE_PT)
```

Use this command to focus on TE port initialization. This example allows two VSANs on the TE port and the port VSAN is 666. Hence the ELP, ESC, and EPP (0x71) go out on VSAN 666. Once the EPP negotiation is complete, we see EFP, DIA, RDI, MR, FSPF, and other updates flow for each allowed VSAN. See Example 58-6.

Example 58-6 Displays SW_ILS Traffic Between Fabric Controllers for all VSANs and Exclude FSPF Hellos and ACK1 Frames.

```
switch(config)# fcan lo bri dis
fc.type==0x22&&((fc.d_id==\"ff.fc.ef\"\|\|fc.s_id==\"ff.fc.ef\"))
Warning:Couldn't obtain netmask info (eth2:no IPv4 address assigned).
Capturing on eth2
9.472181 ff.fc.ef -> ff.fc.61 0x5e0a 0xffff SW_ILS ACA
9.472777 ff.fc.61 -> ff.fc.ef 0x5e0a 0x5e09 SW_ILS SW_ACC (ACA)
9.474551 ff.fc.ef -> ff.fc.61 0x5e0b 0xffff SW_ILS SFC
9.475706 ff.fc.61 -> ff.fc.ef 0x5e0b 0x5e0a SW_ILS SW_ACC (SFC)
9.476694 ff.fc.ef -> ff.fc.61 0x5e0c 0xffff SW_ILS UFC
9.483612 ff.fc.61 -> ff.fc.ef 0x5e0c 0x5e0b SW_ILS SW_ACC (UFC)
9.488187 ff.fc.ef -> ff.fc.61 0x5e0d 0xffff SW_ILS RCA
9.493703 ff.fc.61 -> ff.fc.ef 0x5e0d 0x5e0c SW_ILS SW_ACC (RCA)
```

This example focuses on zone server changes. Prior knowledge of the domain controller ID is required. The switch domain ID where the fcanalyzer is run is x79, the domain controller is FF.FC.79. See Example 58-7.

Example 58-7 Display Switch Internal Link Services (SW_ILS) Traffic to and from Fabric Domain Controller ff.fc.79

```
switch(config)# fcan lo bri dis fc.type==0x22&&((fc.d_id==\ ff.fc.79\
\|\|fc.s_id==\"ff.fc.79\"))
Capturing on eth2
64.053927 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e15 0xffff 0xff -> 0x0 ACA
64.053995 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e16 0xffff 0xff -> 0x0 ACA
64.054599 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e16 0xb1e2 0x5 -> 0xf SW_ACC (ACA)
64.054747 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e15 0x3037 0x4 -> 0xf SW_ACC (ACA)
64.057643 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e17 0xffff 0xff -> 0x0 SFC
64.057696 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e18 0xffff 0xff -> 0x0 SFC
64.058788 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e17 0x3038 0x5 -> 0xf SW_ACC (SFC)
64.059288 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e18 0xb1e3 0x5 -> 0xf SW_ACC (SFC)
64.062011 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e19 0xffff 0xff -> 0x0 UFC
64.062060 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e1a 0xffff 0xff -> 0x0 UFC
64.073513 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e19 0x3039 0x5 -> 0xf SW_ACC (UFC)
64.765306 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e1a 0xb1e4 0x5 -> 0xf SW_ACC (UFC)
64.765572 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e1b 0xffff 0xff -> 0x0 RCA
64.765626 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e1c 0xffff 0xff -> 0x0 RCA
64.766386 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e1b 0x303a 0x4 -> 0xf SW_ACC (RCA)
64.766392 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e1c 0xb1e5 0x5 -> 0xf SW_ACC (RCA)
```

Note You can find the fabric domain controller address in the Mgmt-Id field in the **show fcs ie vsan** command output.

```
switch# show fcs ie vsan 999
IE List for VSAN:999
------------------------------------------------------------------------
IE-WWN                    IE-Type             Mgmt-Id         Mgmt-Addr
------------------------------------------------------------------------
23:e7:00:05:30:00:91:5f   Switch (Remote)     0xfffc04         10.66.78.51
23:e7:00:05:30:00:9b:9f   Switch (Adjacent)   0xfffc01         10.66.78.52
23:e7:00:0d:ec:00:93:81   Switch (Local)      0xfffc79         10.66.78.54
[Total 3 IEs in Fabric]
```

Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters are useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restrict a capture to the specified frames. No other frames are visible until you specify a completely new capture.

The syntax for capture filters is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Ethereal web site (http://www.ethereal.com). Some examples of how you can use this feature follows:

- To capture frames only on a specified VSAN, use this expression:

  ```
  vsan = 1
  ```

- To capture only class F frames, use this expression:

  ```
  class_f
  ```

- To capture only class Fibre Channel ELS frames, use this expression:

  ```
  els
  ```

- To capture only name server frames, use this expression:

  ```
  dns
  ```

- To capture only SCSI command frames, use this expression:

  ```
  fcp_cmd
  ```

Note This feature is part of libpcap and you can obtain more information from http://www.tcpdump.org.

Permitted Capture Filters

This section lists the permitted capture filters.

```
o vsan
o src_port_idx
o dst_port_idx
o sof
o r_ctl
o d_id
o s_id
o type
o seq_id
o seq_cnt
o ox_id
o rx_id
o els
o swils
o fcp_cmd   (FCP Command frames only)
o fcp_data (FCP data frames only)
o fcp_rsp   (FCP response frames only)
o class_f
```

```
o bad_fc
o els_cmd
o swils_cmd
o fcp_lun
o fcp_task_mgmt
o fcp_scsi_cmd
o fcp_status
o gs_type      (Generic Services type)
o gs_subtype   (Generic Services subtype)
o gs_cmd
o gs_reason
o gs_reason_expl
o dns   (name server)
o udns (unzoned name server)
o fcs   (fabric configuration server)
o zs    (zone server)
o fc    (use as fc[x:y] where x is offset and y is length to compare)
o els   (use as els[x:y] similar to fc)
o swils (use as swils[x:y] similar to fc)
o fcp   (use as fcp[x:y] similar to fc)
o fcct (use as fcct[x:y] similar to fc)
```

Loop Monitoring

This section includes the following topics:

About Loop Monitoring

By default, loop monitoring is disabled in all switches in the Cisco MDS 9000 Family. When a disk is removed from a loop port, the loop stays active based on the bypass circuit. Thus the disk removal is not known until you try to communicate with the disk. To detect such removals, the disks can be polled periodically (every 20 seconds).

Caution Changes to the loop monitoring feature should be made by an administrator or individual who is completely familiar with switch operations.

Enabling Loop Monitoring

To enable the loop monitoring feature, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcinterop loop-monitor`	Enables the loop polling for FL ports.
	`switch(config)# no fcinterop loop-monitor`	Disables (default) the loop monitoring feature and reverts the switch to the factory defaults.

Verifying Loop Monitoring Configuration

Use the show running-config command to verify the loop monitoring configuration.

```
switch# show running-config | include loop-monitor
fcinterop loop-monitor
```

The show tech-support Command

The **show tech-support** command is useful when collecting a large amount of information about your switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

The **show tech-support** command displays the output of several **show** commands at once. The output from this command varies depending on your configuration. Use the **show tech-support** command in EXEC mode to display general information about the switch when reporting a problem.

You can choose to have detailed information for each command or even specify the output for a particular interface, module, or VSAN. Each command output is separated by line and the command precedes the output.

Note Explicitly set the **terminal length** command to 0 (zero) to disable auto-scrolling and enable manual scrolling. Use the **show terminal** command to view the configured the terminal size. After obtaining the output of this command, remember to reset your terminal length as required (see the "Setting the Terminal Screen Length" section on page 2-19).

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support** command (see the "Saving Command Output to a File" section on page 2-32). If you save this file, verify you have sufficient space to do so—each of these files may take about 1.8 MB. However, you can zip this file using the **gzip** *filename* command (see the "Compressing and Uncompressing Files" section on page 2-33). Copy the zipped file to the required location using the **copy** command and unzip the file using the **gunzip** command (see the "Copying Files" section on page 2-30).

The default output of the **show tech-support** command includes the output of the following commands:

- **show version**
- **show environment**
- **show module**
- **show hardware**
- **show running-config**
- **show interface**
- **show accounting log**
- **show process**
- **show process log**
- **show processes log details**
- **show flash**

Each command is discussed in both the *Cisco MDS 9000 Family CLI Configuration Guide* and the *Cisco MDS 9000 Family Command Reference*. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide* to obtain debug processes, procedures, and examples.

The show tech-support brief Command

Use the **show tech-support brief** command to obtain a quick, condensed review of your switch configurations. This command provides a summary of the current running state of the switch (see Example 58-8).

The **show tech-support brief** command is useful when collecting information about your switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support brief** command (see the "Saving Command Output to a File" section on page 2-32).

Example 58-8 Displays the Condensed View of Switch Configurations

```
vegas01# show tech-support brief
Switch Name           : vegas01
Switch Type           : DS-X9216-K9-SUP
Kickstart Image       : 1.3(2) bootflash:///m9200-ek9-kickstart-mz.1.3.1.10.bin
System Image          : 1.3(2) bootflash:///m9200-ek9-mz.1.3.1.10.bin
IP Address/Mask       : 10.76.100.164/24
Switch WWN            : 20:00:00:05:30:00:84:9e
No of VSANs           : 9
Configured VSANs      : 1-6,4091-4093

VSAN    1:    name:VSAN0001, state:active, interop mode:default
              domain id:0x6d(109), WWN:20:01:00:05:30:00:84:9f [Principal]
              active-zone:VR, default-zone:deny

VSAN    2:    name:VSAN0002, state:active, interop mode:default
              domain id:0x7d(125), WWN:20:02:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN    3:    name:VSAN0003, state:active, interop mode:default
              domain id:0xbe(190), WWN:20:03:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN    4:    name:VSAN0004, state:active, interop mode:default
              domain id:0x5a(90), WWN:20:04:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN    5:    name:VSAN0005, state:active, interop mode:default
              domain id:0x13(19), WWN:20:05:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN    6:    name:VSAN0006, state:active, interop mode:default
              domain id:0x1f(31), WWN:20:06:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN 4091:    name:VSAN4091, state:active, interop mode:default
              domain id:0x08(8), WWN:2f:fb:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny
```

```
VSAN 4092:    name:VSAN4092, state:active, interop mode:default
              domain id:0x78(120), WWN:2f:fc:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny

VSAN 4093:    name:VSAN4093, state:active, interop mode:default
              domain id:0x77(119), WWN:2f:fd:00:05:30:00:84:9f [Principal]
              active-zone:<NONE>, default-zone:deny


--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          FCOT   Oper  Oper    Port
                  Mode   Trunk                         Mode  Speed   Channel
                         Mode                                 (Gbps)
--------------------------------------------------------------------------------
fc1/1      1      auto   on      fcotAbsent       --     --           --
fc1/2      1      auto   on      fcotAbsent       --     --           --
fc1/3      1      auto   on      fcotAbsent       --     --           --
fc1/4      1      auto   on      fcotAbsent       --     --           --
fc1/5      1      auto   on      notConnected     swl    --           --
fc1/6      1      auto   on      fcotAbsent       --     --           --
fc1/7      1      auto   on      fcotAbsent       --     --           --
fc1/8      1      auto   on      fcotAbsent       --     --           --
fc1/9      1      auto   on      fcotAbsent       --     --           --
fc1/10     1      auto   on      fcotAbsent       --     --           --
fc1/11     1      auto   on      fcotAbsent       --     --           --
fc1/12     1      auto   on      fcotAbsent       --     --           --
fc1/13     1      auto   on      fcotAbsent       --     --           --
fc1/14     1      auto   on      fcotAbsent       --     --           --
fc1/15     1      auto   on      fcotAbsent       --     --           --
fc1/16     1      auto   on      fcotAbsent       --     --           --

--------------------------------------------------------------------------------
Interface          Status                            Speed
                                                     (Gbps)
--------------------------------------------------------------------------------
sup-fc0            up                                1

--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed        MTU
--------------------------------------------------------------------------------
mgmt0                   up         10.76.100.164/24  100 Mbps     1500
```

The show tech-support zone Command

Use the **show tech-support zone** command to obtain information about the zoning configuration on your switch (see Example 58-9).

The output of the **show tech-support zone** command includes the output of the following commands:

- **show zone status vsan**
- **show zone active vsan**
- **show zoneset vsan**
- **show zone vsan**
- **show zone-attribute-group vsan**
- **show zone policy vsan**
- **show zoneset pending active vsan**

- **show zoneset pending vsan**
- **show zone active vsan**
- **show zone pending active vsan**
- **show fcalias pending vsan**
- **show zone-attribute-group pending vsan**
- **show zone policy pending vsan**
- **show zone pending-diff vsan**
- **show zone analysis active vsan**
- **show zone analysis vsan**
- **show zone ess vsan**
- **show zone statistics vsan**
- **show zone statistics lun-zoning vsan**
- **show zone statistics read-only-zoning vsan**

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support zone** command (see the "Saving Command Output to a File" section on page 2-32).

Example 58-9 Displays the Zoning Configurations

```
switch# show tech-support zone vsan 1

`show zone status vsan 1`
VSAN: 1 default-zone: permit distribute: active only Interop: default
mode: basic merge-control: allow session: none
hard-zoning: enabled
Default zone:
qos: disabled broadcast: disabled ronly: disabled
Full Zoning Database :
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Name: vhost-zone Zonesets:1 Zones:9
Status: Activation failed [Error: Unknown error Dom 21]:
at 23:36:44 UTC Dec 19 2005
```

The show tech-support port-channel Command

Use the **show tech-support port-channel** command to obtain information about the PortChannel configuration on your switch (see Example 58-10).

The output of the **show tech-support port-channel** command includes the output of the following commands:

- **show port-channel internal event-history all**
- **show port-channel internal event-history errors**
- **show port-channel internal event-history lock**
- **show port-channel internal mem-stats detail**

- **show port-channel usage**
- **show port-channel summary**
- **show port-channel internal database**
- **show port-channel consistency detail**

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support port-channel** command (see the "Saving Command Output to a File" section on page 2-32).

Example 58-10 Displays the PortChannel Configurations

```
switch# show tech-support port-channel
cp: missing destination file
Try `cp --help' for more information.

`show port-channel internal event-history all`
Low Priority Pending queue: len(0), max len(1) [Wed Jan  4 18:29:18 2006]
High Priority Pending queue: len(0), max len(14) [Wed Jan  4 18:29:18 2006]
PCM Control Block info:
pcm_max_channels        : 128
pcm_max_channel_in_use  : 1
has Vegas Line Card
Total of 1 Vegas Line cards
PCM total_vlans info: 0x0
====================================================
PORT CHANNELS:
====================================================

ALL PORTS:
GigabitEthernet3/1
peer           : 00:00:00:00:00:00:00:00
my wwn         : 00:00:00:00:00:00:00:00
state          : down
update         : none
intent         : unknown
status         : unknown
mode           : on
fcip timeout   : 0 ms
sigloss        : FALSE
flags          :
cfg flags      :
up_time        : 0 usecs after Thu Jan  1 00:00:00 1970
auto pc        : none
auto retry     : 0
last pcp err   : 0 at 0 usecs after Thu Jan  1 00:00:00 1970
No auto create compat failure
...
```

The show tech-support vsan Command

Use the **show tech-support vsan** command to obtain information about the VSAN configuration on your switch (see Example 58-11).

The output of the **show tech-support vsan** command includes the output of the following commands:

- **show vsan**
- **show vsan membership**
- **show interface brief**
- **show port-channel database**
- **show port-channel consistency**
- **show flogi database vsan**
- **show fcdomain vsan**
- **show fcdomain domain-list vsan**
- **show fcdomain address-allocation vsan**
- **show fcns database vsan**
- **show fcs ie vsan**
- **show rscn statistics vsan**
- **show fspf vsan**
- **show fspf database vsan**
- **show span session**
- **show snmp**
- **show zone tech-support vsan**

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support vsan** command (see the "Saving Command Output to a File" section on page 2-32).

Example 58-11 Displays the VSAN Configurations

```
switch# show tech-support vsan 1

`show vsan 1`
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:up


`show vsan 1 membership`
vsan 1 interfaces:
        fc3/1   fc3/2   fc3/3   fc3/4   fc3/5   fc3/6   fc3/7   fc3/8
        fc3/9   fc3/10  fc3/11  fc3/12  fc3/13  fc3/14  port-channel 1  iscsi3/1 iscsi3/2
...
```

The show tech-support fcdomain Command

Use the **show tech-support fcdomain** command to obtain information about the fcdomain configuration on your switch (see Example 58-9).

The output of the **show tech-support fcdomain** command includes the output of the following commands:

- **show fcdomain**
- **show fcdomain domain-list**
- **show fcdomain allowed**
- **show fcdomain pending-diff**
- **show fcdomain address-allocation**
- **show fcdomain address-allocation cache**
- **show fcdomain fcid persistent**
- **show fcdomain internal event-history**
- **show fcdomain internal event-history fcid**
- **show fcdomain internal mem-stats detail**
- **show fcdomain statistics**
- **show fcdomain internal info mts**
- **show fcdomain internal info fcidp-tbl range**

Tip You can save the output of this command to a file by appending **>** (left arrow) and the filename to the **show tech-support fcdomain** command (see the "Saving Command Output to a File" section on page 2-32).

Example 58-12 Displays the fcdomain Configurations

```
switch# show tech-support fcdomain

`show fcdomain status`
fcdomain distribution is disabled

`show fcdomain session-status`

Session parameters for VSAN 1
-----------------------------------
Last Action: none yet
Result: not available

`show fcdomain`

VSAN 1
The local switch is the Principal Switch.

Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:0c:85:90:3e:81
        Running fabric name: 20:01:00:0c:85:90:3e:81
        Running priority: 128
        Current domain ID: 0x72(114)
```

```
Local switch configuration information:
        State: Enabled
        FCID persistence: Enabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)

Principal switch run time information:
        Running priority: 128

No interfaces available.
...
```

IP Network Simulator

The IP Network Simulator tool is supported on the 8-port IP Storage Services (IPS-8) module and 4-port IP Storage Services (IPS-4) module only. You must also have either the SAN extension over IP package for IPS-8 modules (SAN_EXTN_OVER_IP) or SAN extension over IP package for IPS-4 modules (SAN_EXTN_OVER_IP_IPS4) so that you can enable the SAN Extension Tuner, which is a prerequisite for enabling and using the network simulator.

Note You must have a pair of Gigabit Ethernet ports dedicated for each Ethernet path requiring simulation; these ports cannot provide FCIP or iSCSI functionality while simulation occurs. Of course, the remaining ports that are not performing network simulations can run FCIP or iSCSI.

Ports dedicated to network simulation must be adjacent, and always begin with an odd-numbered port. For example, GE 1/1 and GE 1/2 would be a valid pair, while GE 2/2 and GE 2/3 would not.

Network simulator enables you to simulate a variety of IP data network conditions, including the ability to test the impact of network latency. Network simulator is a generic tool that can provide simulation features for all Ethernet traffic; it is not limited to FCIP and iSCSI traffic to or from the Cisco MDS 9000 Family.

The simulation handles full duplex Gigabit Ethernet traffic at full line rate. Figure 58-2 depicts the physical topology using a Cisco MDS 9506 director with an IPS-8 module. GE ports 1 and 2 serve as the network simulator. The FCIP tunnel runs between the Cisco MDS 9506 director port GE 2/1 and the Cisco 9216 module port GE 2/2.

Figure 58-2 Network Simulator - Physical Topology Example

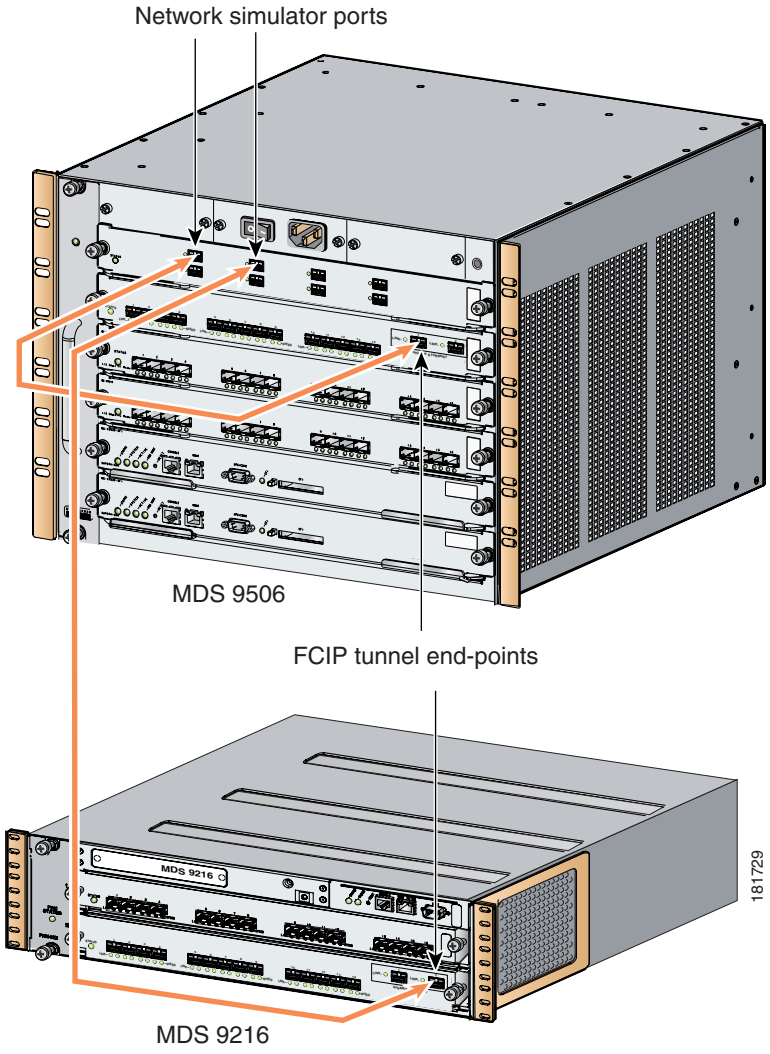

Figure 58-3 depicts the packet flow between the Cisco MDS 9506 and Cisco MDS 9216. Simulations such as delays, drops, and packet reordering are applied independently in each direction. To configure a delay simulation in both directions, you must configure the simulation on both the Cisco MDS 9506 GE 1/1 and 1/2 ports. Simulations are applied to ingress traffic only. All packets received on one Gigabit Ethernet port are sent out of the other Gigabit Ethernet port, and all network configuration simulations are made with respect to the ingress Gigabit Ethernet port.

Figure 58-3 Network Simulator Packet Flow

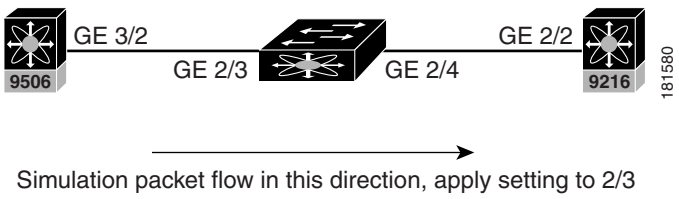

The network simulator tool can simulate the following network functions:

- Network delays (maximum network delays of 150 ms)

- Limiting maximum bandwidth
- Finite queue size
- Dropping packets
- Reordering packets

Enabling the IP Network Simulator

Because the network simulator commands and functionality are part of the SAN Extension Tuner, you must first enable the tuner; after doing so, you can view and use the network simulator commands in EXEC mode.

To enable the network simulator (in this case, on a Cisco MDS 9506 director), follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# san-ext-tuner enable`	Enables the SAN Extension Tuner.
Step 3	`switch(config)# exit` `switch#`	Exits to EXEC mode.
Step 4	`switch# ips netsim enable interface gigabitethernet 2/3 gigabitethernet 2/4`	Configures the pair of Gigabit Ethernet ports in network simulation mode. **Note** The two ports must be adjacent to each; the first port must be an odd-numbered port.
	`switch# ips no netsim enable interface gigabitethernet 2/3 gigabit ethernet 2/4`	Disables network simulation mode and resets the Gigabit Ethernet ports.

Simulating Network Delays

You can configure the network simulator to delay all packets entering the Gigabit Ethernet ports. After configuring the delay in one direction, you need to also enter the same command to introduce the delay in the opposite direction, if desired. You can specify the delay in either milliseconds (allowable range is 0 to150 ms) or microseconds (allowable range is 0 to 150000 µs).

To configure the network simulator to delay all packets entering the Gigabit Ethernet ports 2/3 and 2/4 by 100 ms (round-trip), follow these steps:

	Command	Purpose
Step 1	`switch# ips netsim delay-ms 50 ingress gigabitethernet 2/3`	Configures the network simulator to delay all packets entering the Gigabit Ethernet port 2/3 by 50 ms.
	`switch# ips netsim delay-us 50 ingress gigabitethernet 2/3`	Configures the network simulator to delay all packets entering the Gigabit Ethernet port 2/3 by 50 µs.
Step 2	`switch# ips netsim delay-ms 50 ingress gigabitethernet 2/4`	Configures the network simulator to delay all packets entering the Gigabit Ethernet port 2/4 by 50 ms.
	`switch# ips netsim delay-us 50 ingress gigabitethernet 2/4`	Configures the network simulator to delay all packets entering the Gigabit Ethernet port 2/4 by 50 µs
	`switch# ips netsim delay-ms 0 ingress gigabitethernet 2/3 gigabitethernet 2/4` `switch# ips netsim delay-us 0 ingress gigabitethernet 2/3 gigabitethernet 2/4`	Disables network packet delay simulation.

Simulating Maximum Bandwidth

You can configure the network simulator to restrict the maximum bandwidth in a single direction. Simulating a maximum bandwidth less than that provided by Gigabit Ethernet allows you to control the pacing of packets through the network. So simulating maximum bandwidth in this way actually gives you an idea of the actual bandwidth across a WAN link (for example, an OC3).

You can specify the allowable bandwidth range in either kilobits per second (1000 to 1000000) or megabits per second (1 to 1000).

To configure the network simulator to limit the bandwidth in a specified direction, follow these steps.

	Command	Purpose
Step 1	`switch# ips netsim max-bandwidth-kbps 4500 ingress gigabitethernet 2/3`	Configures the network simulator to limit the bandwidth rate to 4500 kbps for the Gigabit Ethernet port 2/3 in one direction only.
	`switch# ips netsim max-bandwidth-mbps 45 ingress gigabitethernet 2/3`	Configures the network simulator to limit the bandwidth rate to 45 mbps for the Gigabit Ethernet port 2/3 in one direction only.

	Command	Purpose
Step 2	`switch# ips netsim max-bandwidth-kbps 4500 ingress gigabitethernet 2/4`	Configures the network simulator to limit the bandwidth rate to 4500 kbps for the Gigabit Ethernet port 2/4 in one direction only.
	`switch# ips netsim max-bandwidth-mbps 45 ingress gigabitethernet 2/4`	Configures the network simulator to limit the bandwidth rate to 45 mbps for the Gigabit Ethernet port 2/4 in one direction only.

	`switch# ips netsim max-bandwidth-kbps 0 ingress gigabitethernet 2/3`
	
	`switch# ips netsim max-bandwidth-kbps 0 ingress gigabitethernet 2/4`
		Disables network bandwidth rate simulation.

Simulating a Finite Queue Size

You can configure network simulator to simulate a finite queue size in a network device. Data packets are dropped after the queue is full. To simulate a realistic network device, you should specify a queue size of 50 to150 KB. The maximum acceptable queue size is 1000 KB.

To configure the network simulator to simulate a finite queue size, follow these steps.

	Command	Purpose
Step 1	`switch# ips netsim qsize 75 ingress gigabitethernet 2/3`	Configures the network simulator to simulate a finite queue size of 75 KB for the Gigabit Ethernet port 2/3 in one direction only.
Step 2	`switch# ips netsim qsize 75 ingress gigabitethernet 2/4`	Configures the network simulator to simulate a finite queue size of 75 KB for the Gigabit Ethernet port 2/4 in one direction only.
	`switch# ips netsim qsize 1000 ingress gigabitethernet 2/3 gigabitethernet 2/4`	Disables finite queue size simulation.

Simulating Packet Drops

You can configure network simulator to simulate packet drops (even when the queue is not full) randomly (specified as a percentage) or every Nth packet.

Percentage is represented as the number of packets in 10000. For example, if you wish to drop one percent of packets, then you would specify it as 100 packets in 10000. To simulate a realistic scenario for IP networks using random drops, the drop percentage should be between zero and one percent of packet drops in the specified traffic direction.

If you use the optional burst parameter, then the specified number of packets will be dropped each time a decision is made to drop a packet. If you do not specify the burst parameter, then only one packet is dropped each time a decision is made to drop packets. The burst limit for either random or Nth drops is between 1 and 100 packets. Take the burst parameter into account when specifying the percentage of packet drops. For example, if you select random drops of 100 packets in 10,000 (one percent) with a burst size of 2, then 200 packets (or two percent) are dropped every 10,000 packets.

To configure the network simulator to simulate packet drops, follow these steps:

	Command	Purpose
Step 1	`switch# ips netsim drop random 100 burst 1 ingress gigabitethernet 2/3`	Configures the network simulator to simulate random packet drops of 1% for the Gigabit Ethernet port 2/3 in one direction only. The burst is one packet.
	`switch# ips netsim drop nth 100 burst 2 ingress gigabitethernet 2/3`	Configures the network simulator to drop 2 packets after every 100 packets for the Gigabit Ethernet port 2/3 in one direction only (meaning that when the drop is made, two consecutive packages are dropped).
Step 2	`switch# ips netsim drop random 100 burst 1 ingress gigabitethernet 2/4`	Configures the network simulator to simulate a random packet drop of 1% for the Gigabit Ethernet port 2/4 in one direction only. The burst is one packet.
	`switch# ips netsim drop nth 100 burst 2 ingress gigabitethernet 2/4`	Configures the network simulator to drop 2 packets after every 100 packets for the Gigabit Ethernet port 2/4 in one direction only. The burst is two packets, meaning that when the drop is made, two consecutive packages are dropped.
	`switch# ips netsim drop random 0 burst 1 ingress gigabitethernet 2/3` `switch# ips netsim drop nth 0 burst 1 ingress gigabitethernet 2/4`	Disables packet drop simulation.

Simulating Packet Reordering

You can configure network simulator to simulate that a percentage of packets be reordered, either randomly or every Nth packet. Percentage is represented as the number of packets to be reordered in 10000 packets. The acceptable range is between 0 and 10000. So, a specified value of 100 is equal to 1 percent; a value of 1000 is equal to 10 percent.

If you specify the optional distance parameter, then the packet at the head of the queue is reordered with the packet at the distance specified. For example, if you specify a distance of 2 for every 100 packets, then packets 100 and 102 are reordered. The packet sequence would be 1...99, 101, 102, 103...199, 201, 202, 200, 203 and so on. Hence, distance determines how far back in the queue a reordered packet is placed.

To configure the network simulator to simulate packet reordering, follow these steps:

	Command	Purpose
Step 1	`switch# ips netsim reorder random 50 distance 2 ingress gigabitethernet 2/3`	Configures the network simulator to randomly simulate packet reordering at 50% for the Gigabit Ethernet port 2/3 in one direction only. The distance limit is 5.
	`switch# ips netsim reorder nth 50 distance 2 ingress gigabitethernet 2/3`	Configures the network simulator to simulate packet reordering every 50th packet the Gigabit Ethernet port 2/3 in one direction only. The distance limit is 2. So every 50th packet is reordered as the 52nd packet.
Step 2	`switch# ips netsim reorder random 50 distance 2 ingress gigabitethernet 2/4`	Configures the network simulator tool to randomly simulate packet reordering at 50% for the Gigabit Ethernet port 2/4 in one direction only. The distance limit is 2.
	`switch# ips netsim reorder nth 50 distance 2 ingress gigabitethernet 2/4`	Configures the network simulator to simulate packet reordering every 50th packet for the Gigabit Ethernet port 2/4 in one direction only. The distance limit is 2.

	`switch# ips netsim reorder random 0 ingress gigabitethernet 2/3 gigabitethernet 2/4`
	
	`switch# ips netsim reorder nth 0 ingress gigabitethernet 2/3 gigabitethernet 2/4`
		Disables packet reorder simulation.

Displaying IP Network Simulator Statistics

You can view a summary of the IP ports that are currently operating in network simulation mode using the **show ips netsim** command.

```
switch# show ips netsim
Following ports operate in network simulator mode
GigabitEthernet2/3 and GigabitEthernet2/4
```

You can view a summary of the configured parameters and statistics of network simulation using the **show ips stats netsim ingress gigabit ethernet x/y** command. The configuration parameters displayed by default are:

- Delay
- Bandwidth
- Qsize
- Qdelay

The optional configuration parameters are displayed only if they are currently configured on the specified port.

The following network statistics are also displayed:

- Number of packets dropped
- Queue size

- Number of packets reordered
- Average speed

```
switch# show ips stats netsim ingress gigabitethernet 2/3
Network Simulator Configuration for Ingress on GigabitEthernet2/3
    Delay            : 50000 microseconds
    Rate             : 1000000 kbps
    Max_q            : 100000 bytes
    Max_qdelay       : 600000 clocks
    Random Drop %    : 1.00%

Network Simulator Statistics for Ingress on GigabitEthernet2/3
    Dropped (tot)      = 28
    Dropped (netsim)   = 14
    Reordered (netsim) = 0
    Max Qlen(pkt)      = 7
    Qlen    (pkt)      = 0
    Max Qlen (byte)    = 326
    Qlen     (byte)    = 0
    Mintxdel(poll)     = 852
    Mintxdel(ethtx)    = 360
    empty              = 757
    txdel              = 8
    late               = 617
    Average speed      = 0 Kbps
switch# show ips stats netsim ingress gigabitethernet 2/4
Network Simulator Configuration for Ingress on GigabitEthernet2/4
    Delay            : 50000 microseconds
    Rate             : 1000000 kbps
    Max_q            : 100000 bytes
    Max_qdelay       : 600000 clocks
    Reorder nth pkt  : 50
      distance       : 2

Network Simulator Statistics for Ingress on GigabitEthernet2/4
    Dropped (tot)      = 0
    Dropped (netsim)   = 0
    Reordered (netsim) = 2
    Max Qlen(pkt)      = 8
    Qlen    (pkt)      = 0
    Max Qlen (byte)    = 0
    Qlen     (byte)    = 0
    Mintxdel(poll)     = 3788
    Mintxdel(ethtx)    = 360
    empty              = 595
    txdel              = 0
    late               = 335
    Average speed      = 0 Kbps
```

IP Network Simulator Configuration Example

The following example shows how to set up and use the network simulator to introduce a network delay simulation. For continuity, the procedures for creating the Gigabit Ethernet interfaces and enabling the FCIP tunnels are included.

Step 1 Before enabling the network simulator, you must configure two Gigabit Ethernet interfaces to create an FCIP tunnel link (Gigabit Ethernet interfaces 2/3 and 2/4), and then enable the tunnel.

```
switch# config t
switch(config)# interface gigabitethernet 2/3 no shut
```

```
switch(config)# interface gigabitethernet 2/4 no shut
```

Step 2 Enable the SAN Extension Tuner; this is required for the network simulator tool to work.

```
switch(config)# san-ext-tuner enable
switch(config)# exit
```

Step 3 Enable the network simulator on Gigabit Ethernet ports 2/3 and 2/4. Then check that the Gigabit Ethernet ports are operating in network simulation mode.

```
switch# ips netsim enable interface gigabitethernet 2/3 gigabitethernet 2/4
switch# show ips netsim
Following ports operate in network simulator mode
GigabitEthernet2/3 and GigabitEthernet2/4
```

Step 4 Configure a delay of 100 ms round trip (sum of both trips) for all the packets that are arriving on the specified Gigabit Ethernet port.

```
switch# ips netsim delay-ms 50 ingress gigabitethernet 2/3
switch# ips netsim delay-ms 50 ingress gigabitethernet 2/4
```

Step 5 Confirm that the delay you introduced is configured.

```
switch# show ips stats netsim ingress gigabitethernet 2/3
Network Simulator Configuration for Ingress on GigabitEthernet2/3
    Delay            : 50000 microseconds
    Rate             : 1000000 kbps
    Max_q            : 100000 bytes
    Max_qdelay       : 600000 clocks

Network Simulator Statistics for Ingress on GigabitEthernet2/3
    Dropped (tot)    = 0
    Dropped  (ne)    = 0
    Reordered (ne)   = 0
    Max Qlen(pkt)    = 5
    Qlen    (pkt)    = 0
    Max Qlen (byte)  = 0
    Qlen     (byte)  = 0
    Mintxdel(poll)   = 128322
    Mintxdel(ethtx)  = 360
    empty            = 9
    txdel            = 0
    late             = 7
    Average speed    = 0 Kbps
```

Default Settings

Table 58-1 lists the default settings for the features included in this chapter.

Table 58-1 Default Settings for Fabric Troubleshooting Features

Parameters	Default
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

Table 58-1 Default Settings for Fabric Troubleshooting Features (continued)

Parameters	Default
Local capture frame limits	10 frames
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.

CHAPTER 59

Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch. It includes the following sections:

Displaying System Processes

Use the **show processes** command to obtain general information about all processes (see Example 59-1 to Example 59-6).

Example 59-1 Displays System Processes

```
switch# show processes
PID    State  PC        Start_cnt    TTY   Process
-----  -----  --------  -----------  ----  -------------
  868      S  2ae4f33e            1     -  snmpd
  869      S  2acee33e            1     -  rscn
  870      S  2ac36c24            1     -  qos
  871      S  2ac44c24            1     -  port-channel
  872      S  2ac7a33e            1     -  ntp
    -     ER         -            1     -  mdog
    -     NR         -            0     -  vbuilder
```

Where:

- PID = process ID.
- State = process state.
  - D = uninterruptible sleep (usually I/O).
  - R = runnable (on run queue).
  - S = sleeping.
  - T = traced or stopped.
  - Z = defunct (“zombie”) process.
- NR = not running.
- ER = should be running but currently not-running.
- PC = current program counter in hex format.
- Start_cnt = number of times a process has been started (or restarted).
- TTY = terminal that controls the process. A hyphen usually means a daemon not running on any particular TTY.
- Process = name of the process.

Example 59-2 Displays CPU Utilization Information

```
switch# show processes cpu
PID    Runtime(ms)  Invoked   uSecs  1Sec   Process
-----  -----------  --------  -----  -----  -----------
  842         3807    137001     27    0.0  sysmgr
 1112         1220     67974     17    0.0  syslogd
 1269          220     13568     16    0.0  fcfwd
 1276         2901     15419    188    0.0  zone
 1277          738     21010     35    0.0  xbar_client
 1278         1159      6789    170    0.0  wwn
 1279          515     67617      7    0.0  vsan
```

Where:

- Runtime (ms) = CPU time the process has used, expressed in milliseconds.
- Invoked = number of times the process has been invoked.
- uSecs = microseconds of CPU time on average for each process invocation.
- 1Sec = CPU utilization in percentage for the last one second.

Example 59-3 Displays Process Log Information

```
switch# show processes log
Process           PID     Normal-exit  Stack-trace  Core     Log-create-time
----------------  ------  -----------  -----------  -------  ---------------
fspf              1339              N            Y        N  Jan  5 04:25
lcm               1559              N            Y        N  Jan  2 04:49
rib               1741              N            Y        N  Jan  1 06:05
```

Where:

- Normal-exit = whether or not the process exited normally.
- Stack-trace = whether or not there is a stack trace in the log.

- Core = whether or not there exists a core file.
- Log-create-time = when the log file got generated.

Example 59-4 Displays Detail Log Information About a Process

```
switch# show processes log pid 1339
Service: fspf
Description: FSPF Routing Protocol Application

Started at Sat Jan  5 03:23:44 1980 (545631 us)
Stopped at Sat Jan  5 04:25:57 1980 (819598 us)
Uptime: 1 hours 2 minutes 2 seconds

Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 9 (no core)
CWD: /var/sysmgr/work

Virtual Memory:

    CODE      08048000 - 0809A100
    DATA      0809B100 - 0809B65C
    BRK       0809D988 - 080CD000
    STACK     7FFFFD20
    TOTAL     23764 KB

Register Set:

    EBX 00000005         ECX 7FFFF8CC         EDX 00000000
    ESI 00000000         EDI 7FFFF6CC         EBP 7FFFF95C
    EAX FFFFFDFE         XDS 8010002B         XES 0000002B
    EAX 0000008E (orig)  EIP 2ACE133E         XCS 00000023
    EFL 00000207         ESP 7FFFF654         XSS 0000002B

Stack: 1740 bytes. ESP 7FFFF654, TOP 7FFFFD20

0x7FFFF654: 00000000 00000008 00000003 08051E95 ................
0x7FFFF664: 00000005 7FFFF8CC 00000000 00000000 ................
0x7FFFF674: 7FFFF6CC 00000001 7FFFF95C 080522CD ........\...."..
0x7FFFF684: 7FFFF9A4 00000008 7FFFFC34 2AC1F18C ........4......*
```

Example 59-5 Displays All Process Log Details

```
switch# show processes log details
================================================
Service: snmpd
Description: SNMP Agent

Started at Wed Jan  9 00:14:55 1980 (597263 us)
Stopped at Fri Jan 11 10:08:36 1980 (649860 us)
Uptime: 2 days 9 hours 53 minutes 53 seconds

Start type: SRV_OPTION_RESTART_STATEFUL (24)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 6 (core dumped)
CWD: /var/sysmgr/work

Virtual Memory:

    CODE      08048000 - 0804C4A0
    DATA      0804D4A0 - 0804D770
    BRK       0804DFC4 - 0818F000
```

```
    STACK     7FFFFCE0
    TOTAL     26656 KB
...
```

Example 59-6 Displays Memory Information About Processes

```
switch# show processes memory
PID    MemAlloc  StackBase/Ptr      Process
-----  --------  -----------------  ---------------
 1277    120632  7ffffcd0/7fffefe4  xbar_client
 1278     56800  7ffffce0/7ffffb5c  wwn
 1279   1210220  7ffffce0/7ffffbac  vsan
 1293    386144  7ffffcf0/7fffebd4  span
 1294   1396892  7ffffce0/7fffdff4  snmpd
 1295    214528  7ffffcf0/7ffff904  rscn
 1296     42064  7ffffce0/7ffffb5c  qos
```

Where:

- MemAlloc = total memory allocated by the process.
- StackBase/Ptr = process stack base and current stack pointer in hex format.

Displaying System Status

Use the **show system** command to display system-related status information (see Example 59-7 to Example 59-10.

Example 59-7 Displays Default Switch Port States

```
switch# show system default switchport
System default port state is down
System default trunk mode is on
```

Example 59-8 Displays Error Information for a Specified ID

```
switch# show system error-id 0x401D0019
Error Facility: module
Error Description: Failed to stop Linecard Async Notification.
```

Example 59-9 Displays the System Reset Information

```
switch# Show system reset-reason module 5
----- reset reason for module 5 -----
1) At 224801 usecs after Fri Nov 21 16:36:40 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
2) At 922828 usecs after Fri Nov 21 16:02:48 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
3) At 318034 usecs after Fri Nov 21 14:03:36 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
```

```
4) At 255842 usecs after Wed Nov 19 00:07:49 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
```

The **show system reset-reason** command displays the following information:

- In a Cisco MDS 9513 Director, the last four reset-reason codes for the supervisor module in slot 7 and slot 8 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9506 or Cisco MDS 9509 switch, the last four reset-reason codes for the supervisor module in slot 5 and slot 6 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9200 Series switch, the last four reset-reason codes for the supervisor module in slot 1 are displayed.
- The **show system reset-reason module** *number* command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module are not displayed.

Use the **clear system reset-reason** command to clear the reset-reason information stored in NVRAM and volatile persistent storage.

- In a Cisco MDS 9500 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active and standby supervisor modules.
- In a Cisco MDS 9200 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active supervisor module.

Example 59-10 Displays System Uptime

```
switch# show system uptime
Start Time: Sun Oct 13 18:09:23 2030
Up Time:    0 days, 9 hours, 46 minutes, 26 seconds
```

Use the **show system resources** command to display system-related CPU and memory statistics (see Example 59-11).

Example 59-11 Displays System-Related CPU and Memory Information

```
switch# show system resources
Load average:   1 minute: 0.43   5 minutes: 0.17   15 minutes: 0.11
Processes   :   100 total, 2 running
CPU states  :   0.0% user,   0.0% kernel,   100.0% idle
Memory usage:   1027628K total,    313424K used,    714204K free
                   3620K buffers,   22278K cache
```

Where:

- Load average—Displays the number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes—Displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states—Displays the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage—Displays the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the *used* memory statistics.

Core and Log Files

This section the following topics:

Displaying Core Status

Use the **show system cores** command to display the currently configured scheme for copying cores. See Examples 59-12 to 59-14.

Example 59-12 Displays the Status of System Cores

```
switch# show system cores
Transfer of cores is enabled
```

Example 59-13 Displays All Cores Available for Upload from the Active Supervisor Module

```
switch# show cores
Module-num  Process-name    PID     Core-create-time
----------  ------------    ---     ----------------
5           fspf            1524    Nov 9 03:11
6           fcc             919     Nov 9 03:09
8           acltcam         285     Nov 9 03:09
8           fib             283     Nov 9 03:08
```

Where `Module-num` shows the slot number on which the core was generated. In this example, the `fspf core` was generated on the active supervisor module (slot 5), `fcc` was generated on the standby supervisor module (slot 6), and `acltcam` and `fib` were generated on the switching module (slot 8).

Example 59-14 Displays Logs on the Local System

```
switch# show processes log
Process           PID     Normal-exit  Stack  Core   Log-create-time
----------------  ------  -----------  -----  -----  ---------------
ExceptionLog      2862              N      Y      N  Wed Aug  6 15:08:34 2003
acl               2299              N      Y      N  Tue Oct 28 02:50:01 2003
bios_daemon       2227              N      Y      N  Mon Sep 29 15:30:51 2003
capability        2373              N      Y      N  Tue Aug 19 13:30:02 2003
core-client       2262              N      Y      N  Mon Sep 29 15:30:51 2003
fcanalyzer        5623              N      Y      N  Fri Sep 26 20:45:09 2003
fcd               12996             N      Y      N  Fri Oct 17 20:35:01 2003
fcdomain          2410              N      Y      N  Thu Jun 12 09:30:58 2003
ficon             2708              N      Y      N  Wed Nov 12 18:34:02 2003
ficonstat         9640              N      Y      N  Tue Sep 30 22:55:03 2003
flogi             1300              N      Y      N  Fri Jun 20 08:52:33 2003
idehsd            2176              N      Y      N  Tue Jun 24 05:10:56 2003
lmgrd             2220              N      N      N  Mon Sep 29 15:30:51 2003
```

```
platform          2840          N      Y      N  Sat Oct 11 18:29:42 2003
port-security     3098          N      Y      N  Sun Sep 14 22:10:28 2003
port              11818         N      Y      N  Mon Nov 17 23:13:37 2003
rlir              3195          N      Y      N  Fri Jun 27 18:01:05 2003
rscn              2319          N      Y      N  Mon Sep 29 21:19:14 2003
securityd         2239          N      N      N  Thu Oct 16 18:51:39 2003
snmpd             2364          N      Y      N  Mon Nov 17 23:19:39 2003
span              2220          N      Y      N  Mon Sep 29 21:19:13 2003
syslogd           2076          N      Y      N  Sat Oct 11 18:29:40 2003
tcap              2864          N      Y      N  Wed Aug  6 15:09:04 2003
tftpd             2021          N      Y      N  Mon Sep 29 15:30:51 2003
vpm               2930          N      N      N  Mon Nov 17 19:14:33 2003
```

Saving Cores

You can save cores (from the active supervisor module, the standby supervisor module, or any switching module) to an external CompactFlash (slot 0) or to a TFTP server in one of two ways:

- On demand—Copies a single file based on the provided process ID.
- Periodically—Copies core files periodically as configured by the user.

A new scheme overwrites any previously issued scheme. For example, if you perform another core log copy task, the cores are periodically saved to the new location or file.

Tip Be sure to create any required directory before performing this task. If the directory specified by this task does not exist, the switch software logs a system message each time a copy cores is attempted.

To copy the core and log files on demand, follow this step:

	Command	Purpose
Step 1	`switch# show cores`	
Step 2	`switch# copy core:7407 slot0:coreSample`	Copies the core file with the process ID 7407 as coreSample in slot 0.
	`switch# copy core://5/1524 tftp:/1.1.1.1/abcd`	Copies cores (if any) of a process with PID 1524 generated on slot 5[1] or slot 7[2] to the TFTP server at IPv4 address 1.1.1.1. **Note** You can also use IPv6 addresses to identify the TFTP server.

1. Cisco MDS 9506 or Cisco MDS 9509 switch
2. Cisco MDS 9513 Director

- If the core file for the specified process ID is not available, you see the following response:

```
switch# copy core:133 slot0:foo
No core file found with pid 133
```

- If two core files exist with the same process ID, only one file is copied:

```
switch# copy core:7407 slot0:foo1
2 core files found with pid 7407
Only "/isan/tmp/logs/calc_server_log.7407.tar.gz" will be copied to the destination.
```

To copy the core and log files periodically, follow these steps:

	Command	Purpose
Step 1	`switch# show system cores`	
Step 2	`switch# config t`	Enters configuration mode.
Step 3	`switch(config)# system cores slot0:coreSample`	Copies the core file (coreSample) to slot 0.
	`switch(config)# system cores tftp:/1.1.1.1/abcd`	Copies the core file (abcd) in the specified directory on the TFTP server at IPv4 address 1.1.1.1. **Note** You can also use IPv6 addresses to identify the TFTP server.
	`switch(config)# no system cores`	Disables the core files copying feature.

Saving the Last Core to CompactFlash

This last core dump is automatically saved to CompactFlash in the /mnt/pss/ partition before the switchover or reboot occurs. Three minutes after the supervisor module reboots, the saved last core is restored from the Flash partition (/mnt/pss) back to its original RAM location. This restoration is a background process and is not visible to the user.

Tip The timestamp on the restored last core file displays the time when the supervisor booted up—not when the last core was actually dumped. To obtain the exact time of the last core dump, check the corresponding log file with the same PID.

To view the last core information, issue the **show cores** command in EXEC mode.

To view the time of the actual last core dump, issue the **show process log** command in EXEC mode.

Clearing the Core Directory

Use the **clear cores** command to clean out the core directory. The software keeps the last few cores per service and per slot and clears all other cores present on the active supervisor module.

```
switch# clear cores
```

Kernel Core Dumps

Caution Changes to the kernel cores should be made by an administrator or individual who is completely familiar with switch operations.

When a specific module's operating system (OS) crashes, it is sometimes useful to obtain a full copy of the memory image (called a *kernel core dump*) to identify the cause of the crash. When the module experiences a kernel core dump it triggers the proxy server configured on the supervisor. The supervisor

sends the module's OS kernel core dump to the Cisco MDS 9000 System Debug Server. Similarly, if the supervisor OS fails, the supervisor sends its OS kernel core dump to the Cisco MDS 9000 System Debug Server.

Note The Cisco MDS 9000 System Debug Server is a Cisco application that runs on Linux. It creates a repository for kernel core dumps. You can download the Cisco MDS 9000 System Debug Server from the Cisco.com website at http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml.

Kernel core dumps are only useful to your technical support representative. The kernel core dump file, which is a large binary file, must be transferred to an external server that resides on the same physical LAN as the switch. The core dump is subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

Tip Core dumps take up disk space on the Cisco MDS 9000 System Debug Server application. If all levels of core dumps (**level all** option) are configured, you need to ensure that a minimum of 1 GB of disk space is available on the Linux server running the Cisco MDS 9000 System Debug Server application to accept the dump. If the process does not have sufficient space to complete the generation, the module resets itself. All changes made to kernel cores are saved to the running configuration.

This section includes the following topics:

Configuring External Servers

To configure the external server using IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# kernel core target 10.50.5.5` `succeeded`	Configures the external server's IPv4 address. **Note** IPv6 addresses are not supported for kernel core targets.

Configuring Module Parameters

To configure the module parameters, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# kernel core module 5` `succeeded`	Configures kernel core generation for module 5.
	`switch(config)# kernel core module 5 level header` `succeeded`	Configures kernel core generation for module 5, and limits the generation to header-level cores.
Step 3	`switch(config)# kernel core limit 2` `succeeded`	Configures kernel core generations for two modules. The default is 1 module.

Displaying Kernel Core Information

All changes made to the kernel cores may be viewed using the **show running-config** command. Alternatively, use the **show kernel cores** command to view specific configuration changes (see Example 59-15 to Example 59-17).

Example 59-15 Displays the Core Limit

```
switch# show kernel core limit
2
```

Example 59-16 Displays the External Server

```
switch# show kernel core target
10.50.5.5
```

Example 59-17 Displays the Core Settings for the Specified Module

```
switch# show kernel core module 5
module 5 core is enabled
         level is header
         dst_ip is 10.50.5.5
         src_port is 6671
         dst_port is 6666
         dump_dev_name is eth1
         dst_mac_addr is 00:00:0C:07:AC:01
```

Online System Health Management

The Online Health Management System (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

About Online System Health Management

The Online Health Management System (OHMS) is a hardware fault detection and recovery feature. It runs on all Cisco MDS switching, services, and supervisor modules and ensures the general health of any switch in the Cisco MDS 9000 Family. The OHMS monitors system hardware in the following ways:

- The OHMS component running on the active supervisor maintains control over all other OHMS components running on the other modules in the switch.
- The system health application running in the standby supervisor module only monitors the standby supervisor module—if that module is available in the HA standby mode. See the "HA Switchover Characteristics" section on page 9-2.

The OHMS application launches a daemon process in all modules and runs multiple tests on each module to test individual module components. The tests run at preconfigured intervals, cover all major fault points, and isolate any failing component in the MDS switch. The OHMS running on the active supervisor maintains control over all other OHMS components running on all other modules in the switch.

On detecting a fault, the system health application attempts the following recovery actions:

- Performs additional testing to isolate the faulty component
- Attempts to reconfigure the component by retrieving its configuration information from persistent storage.
- If unable to recover, sends Call Home notifications, system messages and exception logs; and shuts down and discontinues testing the failed module or component (such as an interface)
- Sends Call Home and system messages and exception logs as soon as it detects a failure.
- Shuts down the failing module or component (such as an interface).
- Isolates failed ports from further testing.
- Reports the failure to the appropriate software component.
- Switches to the standby supervisor module, if an error is detected on the active supervisor module and a standby supervisor module exists in the Cisco MDS switch. After the switchover, the new active supervisor module restarts the active supervisor tests.
- Reloads the switch if a standby supervisor module does not exist in the switch.
- Provides CLI support to view, test, and obtain test run statistics or change the system health test configuration on the switch.
- Performs tests to focus on the problem area.

Each module is configured to run the test relevant to that module. You can change the default parameters of the test in each module as required.

System Health Initiation

By default, the system health feature is enabled in each switch in the Cisco MDS 9000 Family.

To disable or enable this feature in any switch in the Cisco MDS 9000 Family, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no system health` `System Health is disabled.`	Disables system health from running tests in this switch.
	`switch(config)# system health` `System Health is enabled.`	Enables (default) system health to run tests in this switch.
Step 3	`switch(config)# no system health interface fc8/1` `System health for interface fc8/13 is disabled.`	Disables system health from testing the specified interface.
	`switch(config)# system health interface fc8/1` `System health for interface fc8/13 is enabled.`	Enables (default) system health to test for the specified interface.

Loopback Test Configuration Frequency

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured frequency—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run at frequencies ranging from 5 seconds (default) to 255 seconds. If you do not configure the loopback frequency value, the default frequency of 5 seconds is used for all modules in the switch. Loopback test frequencies can be altered for each module.

To configure the frequency of loopback tests for all modules on a switch, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# system health loopback frequency 50` `The new frequency is set at 50 Seconds.`	Configures the loopback frequency to 50 seconds. The default loopback frequency is 5 seconds. The valid range is from 5 to 255 seconds.

Loopback Test Configuration Frame Length

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured size—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run with frame sizes ranging from 0 bytes to 128 bytes. If you do not configure the loopback frame length value, the switch generates random frame lengths for all modules in the switch (auto mode). Loopback test frame lengths can be altered for each module.

To configure the frame length for loopback tests for all modules on a switch, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# system health loopback frame-length 128`	Configures the loopback frame length to 128 bytes. The valid range is 0 to 128 bytes.
Step 3	`switch(config)# system health loopback frame-length auto`	Configures the loopback frame length to automatically generate random lengths (default).

To verify the loopback frequency configuration, use the **show system health loopback frame-length** command.

```
switch# show system health loopback frame-length
Loopback frame length is set to auto-size between 0-128 bytes
```

Hardware Failure Action

The failure-action command controls the Cisco SAN-OS software from taking any action if a hardware failure is determined while running the tests.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family—action is taken if a failure is determined and the failed component is isolated from further testing.

Failure action is controlled at individual test levels (per module), at the module level (for all tests), or for the entire switch.

To configure failure action in a switch, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# system health failure-action` `System health global failure action is now enabled.`	Enables the switch to take failure action (default).
Step 3	`switch(config)# no system health failure-action` `System health global failure action now disabled.`	Reverts the switch configuration to prevent failure action being taken.
Step 4	`switch(config)# system health module 1 failure-action` `System health failure action for module 1 is now enabled.`	Enables switch to take failure action for failures in module 1.
Step 5	`switch(config)# no system health module 1 loopback failure-action` `System health failure action for module 1 loopback test is now disabled.`	Prevents the switch from taking action on failures determined by the loopback test in module 1.

Test Run Requirements

Enabling a test does not guarantee that a test will run.

Tests on a given interface or module only run if you enable system health for all of the following items:

- The entire switch.
- The required module.
- The required interface.

Tip The test will not run if system health is disabled in any combination. If system health is disabled to run tests, the test status shows up as disabled.

Tip If the specific module or interface is enabled to run tests, but is not running the tests due to system health being disabled, then tests show up as enabled (not running).

Tests for a Specified Module

The system health feature in the SAN-OS software performs tests in the following areas:

- Active supervisor's in-band connectivity to the fabric.
- Standby supervisor's arbiter availability.
- Bootflash connectivity and accessibility on all modules.
- EOBC connectivity and accessibility on all modules.
- Data path integrity for each interface on all modules.
- Management port's connectivity.
- Caching Services Module (CSM) batteries (for temperature, age, full-charge capacity, (dis)charge ability and backup capability) and cache disks (for connectivity, accessibility and raw disk I/O).
- User-driven test for external connectivity verification, port is shut down during the test (Fibre Channel ports only).
- User-driven test for internal connectivity verification (Fibre Channel and iSCSI ports).

To perform the required test on a specific module, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
	Note The following steps can be performed in any order.	
Step 2	`switch(config)# system health module 8 battery-charger` `battery-charger test is not configured to run on module 8.`	Enables the battery-charger test on both batteries in the CSM residing in slot 8. If the switch does not have a CSM in slot 8, this message is issued.

	Command	Purpose
Step 3	`switch(config)# system health module 8 cache-disk` `cache-disk test is not configured to run on module 8.`	Enables the cache-disk test on both disks in the CSM residing in slot 8. If the switch does not have a CSM in slot 8, this message is issued.
	Note The various options for each test are described in the next step. Each command can be configured in any order. The various options are presented in the same step for documentation purposes.	
Step 4	`switch(config)# system health module 8 bootflash` `System health for module 8 Bootflash is already enabled.`	Enables the bootflash test on module in slot 8.
	`switch(config)# system health module 8 bootflash frequency 200` `The new frequency is set at 200 Seconds.`	Sets the new frequency of the bootflash test on module 8 to 200 seconds.
Step 5	`switch(config)# system health module 8 eobc` `System health for module 8 EOBC is now enabled.`	Enables the EOBC test on module in slot 8.
Step 6	`switch(config)# system health module 8 loopback` `System health for module 8 EOBC is now enabled.`	Enables the loopback test on module in slot 8.
Step 7	`switch(config)# system health module 5 management` `System health for module 8 EOBC is now enabled.`	Enables the management test on module in slot 5.

Clearing Previous Error Reports

You can clear the error history for Fibre Channel interfaces, iSCSI interfaces, an entire module, or one particular test for an entire module. By clearing the history, you are directing the software to retest all failed components that were previously excluded from tests.

If you previously enabled the failure-action option for a period of time (for example, one week) to prevent OHMS from taking any action when a failure is encountered and after that week you are now ready to start receiving these errors again, then you must clear the system health error status for each test.

Tip The management port test cannot be run on a standby supervisor module.

Use the EXEC-level **system health clear-errors** command at the interface or module level to erase any previous error conditions logged by the system health application. The **battery-charger**, the **bootflash**, the **cache-disk**, the **eobc**, the **inband**, the **loopback**, and the **mgmt** test options can be individually specified for a given module.

The following example clears the error history for the specified Fibre Channel interface:

```
switch# system health clear-errors interface fc 3/1
```

The following example clears the error history for the specified module:

```
switch# system health clear-errors module 3
```

The following example clears the management test error history for the specified module:

```
switch# system health clear-errors module 1 mgmt
```

Performing Internal Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. Internal loopback tests send and receive FC2 frames to/from the same ports and provide the round trip time taken in microseconds. These tests are available for Fibre Channel, IPS, and iSCSI interfaces.

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module.

```
switch# system health internal-loopback interface iscsi 8/1
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame count configured on the switch.

```
switch# system health internal-loopback interface iscsi 8/1 frame-count 20
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame length configured on the switch.

```
switch# system health internal-loopback interface iscsi 8/1 frame-count 32
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error:

```
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback,
analysis complete Failed device ID 3 on module 1
```

Performing External Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. External loopback tests send and receive FC2 frames to/from the same port or between two ports.

You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. If you are testing to/from the same port, you need a special loop cable. If you are testing to/from different ports, you can use a regular cable. This test is only available for Fibre Channel interfaces.

Use the EXEC-level **system health external-loopback interface** *interface* command to run this test on demand for external devices connected to a switch that is part of a long-haul network.

```
switch# system health external-loopback interface fc 3/1
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback source** *interface* **destination interface** *interface* command to run this test on demand between two ports on the switch.

```
switch# system health external-loopback source interface fc 3/1 destination interface fc
3/2
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 and interface fc3/2 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback** *interface* **frame-count** command to run this test on demand for external devices connected to a switch that is part of a long-haul network and override the frame count configured on the switch.

```
switch# system health external-loopback interface fc 3/1 frame-count 10
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback** *interface* **frame-length** command to run this test on demand for external devices connected to a switch that is part of a long-haul network and override the frame length configured on the switch.

```
switch# system health external-loopback interface fc 3/1 frame-length 64
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the **system health external-loopback** *interface* **force** command to shut down the required interface directly without a back out confirmation.

```
switch# system health external-loopback interface fc 3/1 force
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error:

```
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback,
analysis complete Failed device ID 3 on module 1
```

Performing Serdes Loopbacks

Serializer/Deserializer (serdes) loopback tests the hardware for a port. These tests are available for Fibre Channel interfaces.

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module.

```
switch# system health serdes-loopback interface fc 3/1
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame count configured on the switch.

```
switch# system health serdes-loopback interface fc 3/1 frame-count 10
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame length configured on the switch.

```
switch# system health serdes-loopback interface fc 3/1 frame-length 32
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error: `External loopback test on interface fc 3/1 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 3`

Interpreting the Current Status

The status of each module or test depends on the current configured state of the OHMS test in that particular module (see Table 59-1).

Table 59-1 OHMS Configured Status for Tests and Modules

Status	Description
Enabled	You have currently enabled the test in this module and the test is not running.
Disabled	You have currently disabled the test in this module.
Running	You have enabled the test and the test is currently running in this module.
Failing	This state is displayed if a failure is imminent for the test running in this module—possibility of test recovery exists in this state.
Failed	The test has failed in this module—and the state cannot be recovered.
Stopped	The test has been internally stopped in this module by the Cisco SAN-OS software.
Internal failure	The test encountered an internal failure in this module. For example, the system health application is not able to open a socket as part of the test procedure.
Diags failed	The startup diagnostics has failed for this module or interface.
On demand	The system health external-loopback or the system health internal-loopback tests are currently running in this module. Only these two commands can be issued on demand.
Suspended	Only encountered in the MDS 9100 Series due to one oversubscribed port moving to a E or TE port mode. If one oversubscribed port moves to this mode, the other three oversubscribed ports in the group are suspended.

The status of each test in each module is visible when you display any of the **show system health** commands. See the "Displaying System Health" section on page 59-18.

Displaying System Health

Use the **show system health** command to display system-related status information (see Example 59-18 to Example 59-23).

Example 59-18 Displays the Current Health of All Modules in the Switch

```
switch# show system health
```

```
Current health information for module 2.

Test                       Frequency    Status          Action
------------------------------------------------------------------
Bootflash                  5 Sec        Running         Enabled
EOBC                       5 Sec        Running         Enabled
Loopback                   5 Sec        Running         Enabled
------------------------------------------------------------------

Current health information for module 6.

Test                       Frequency    Status          Action
------------------------------------------------------------------
InBand                     5 Sec        Running         Enabled
Bootflash                  5 Sec        Running         Enabled
EOBC                       5 Sec        Running         Enabled
Management Port            5 Sec        Running         Enabled
------------------------------------------------------------------
```

Example 59-19 Displays the Current Health of a Specified Module

```
switch# show system health module 8

Current health information for module 8.

Test                       Frequency    Status          Action
------------------------------------------------------------------
Bootflash                  5 Sec        Running         Enabled
EOBC                       5 Sec        Running         Enabled
Loopback                   5 Sec        Running         Enabled
------------------------------------------------------------------
```

Example 59-20 Displays Health Statistics for All Modules

```
switch# show system health statistics

Test statistics for module # 1
-------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
Bootflash           Running             5s    12900   12900       0     0    0
EOBC                Running             5s    12900   12900       0     0    0
Loopback            Running             5s    12900   12900       0     0    0
-------------------------------------------------------------------------------

Test statistics for module # 3
-------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
Bootflash           Running             5s    12890   12890       0     0    0
EOBC                Running             5s    12890   12890       0     0    0
Loopback            Running             5s    12892   12892       0     0    0
-------------------------------------------------------------------------------

Test statistics for module # 5
-------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
-------------------------------------------------------------------------------
InBand              Running             5s    12911   12911       0     0    0
```

```
Bootflash           Running             5s    12911   12911       0     0    0
EOBC                Running             5s    12911   12911       0     0    0
Management Port     Running             5s    12911   12911       0     0    0
------------------------------------------------------------------------------

Test statistics for module # 6
------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
InBand              Running             5s    12907   12907       0     0    0
Bootflash           Running             5s    12907   12907       0     0    0
EOBC                Running             5s    12907   12907       0     0    0
------------------------------------------------------------------------------

Test statistics for module # 8
------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12895   12895       0     0    0
EOBC                Running             5s    12895   12895       0     0    0
Loopback            Running             5s    12896   12896       0     0    0
------------------------------------------------------------------------------
```

Example 59-21 Displays Statistics for a Specified Module

```
switch# show system health statistics module 3

Test statistics for module # 3
------------------------------------------------------------------------------
Test Name           State            Freq(s)    Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12932   12932       0     0    0
EOBC                Running             5s    12932   12932       0     0    0
Loopback            Running             5s    12934   12934       0     0    0
------------------------------------------------------------------------------
```

Example 59-22 Displays Loopback Test Statistics for the Entire Switch

```
switch# show system health statistics loopback
-----------------------------------------------------------------
Mod Port Status               Run      Pass      Fail   CFail Errs
  1   16 Running             12953    12953        0       0    0
  3   32 Running             12945    12945        0       0    0
  8    8 Running             12949    12949        0       0    0
-----------------------------------------------------------------
```

Example 59-23 Displays Loopback Test Statistics for a Specified Interface

```
switch# show system health statistics loopback interface fc 3/1
-----------------------------------------------------------------
Mod Port Status               Run      Pass      Fail   CFail Errs
  3    1 Running                0         0         0       0    0
-----------------------------------------------------------------
```

Note Interface-specific counters will remain at zero unless the module-specific loopback test reports errors or failures.

Example 59-24 Displays the Loopback Test Time Log for All Modules

```
switch# show system health statistics loopback timelog
-----------------------------------------------------------------
Mod        Samples     Min(usecs)     Max(usecs)     Ave(usecs)
  1           1872            149            364            222
  3           1862            415            743            549
  8           1865            134            455            349
-----------------------------------------------------------------
```

Example 59-25 Displays the Loopback Test Time Log for a Specified Module

```
switch# show system health statistics loopback module 8 timelog
-----------------------------------------------------------------
Mod        Samples     Min(usecs)     Max(usecs)     Ave(usecs)
  8           1867            134            455            349
-----------------------------------------------------------------
```

On-Board Failure Logging

The Generation 2 Fibre Channel switching modules provide the facility to log failure data to persistent storage, which can be retrieved and displayed for analysis. This on-board failure logging (OBFL) feature stores failure and environmental information in nonvolatile memory on the module. The information will help in post-mortem analysis of failed cards.

This section includes the following topics:

About OBFL

OBFL data is stored in the existing CompactFlash on the module. OBFL uses the persistent logging (PLOG) facility available in the module firmware to store data in the CompactFlash. It also provides the mechanism to retrieve the stored data.

The data stored by the OBFL facility includes the following:

- Time of initial power-on
- Slot number of the card in the chassis
- Initial temperature of the card
- Firmware, BIOS, FPGA, and ASIC versions
- Serial number of the card
- Stack trace for crashes
- CPU hog information
- Memory leak information

- Software error messages
- Hardware exception logs
- Environmental history
- OBFL specific history information
- ASIC interrupt and error statistics history
- ASIC register dumps

Configuring OBFL for the Switch

To configure OBFL for all the modules on the switch, follow these steps

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# hw-module logging onboard`	Enables all OBFL features.
	`switch(config)# hw-module logging onboard cpu-hog`	Enables the OBFL CPU hog events.
	`switch(config)# hw-module logging onboard environmental-history`	Enables the OBFL environmental history.
	`switch(config)# hw-module logging onboard error-stats`	Enables the OBFL error statistics.
	`switch(config)# hw-module logging onboard interrupt-stats`	Enables the OBFL interrupt statistics.
	`switch(config)# hw-module logging onboard mem-leak`	Enables the OBFL memory leak events.
	`switch(config)# hw-module logging onboard miscellaneous-error`	Enables the OBFL miscellaneous information.
	`switch(config)# hw-module logging onboard obfl-log`	Enables the boot uptime, device version, and OBFL history.
	`switch(config)# no hw-module logging onboard`	Disables all OBFL features.

Use the **show logging onboard status** command to display the configuration status of OBFL.

```
switch# show logging onboard status

    Switch OBFL Log:                                        Enabled

    Module:  6 OBFL Log:                                    Enabled
    error-stats                                             Enabled
    exception-log                                           Enabled
    miscellaneous-error                                     Enabled
    obfl-log (boot-uptime/device-version/obfl-history)      Enabled
    system-health                                           Enabled
    stack-trace                                             Enabled
```

Configuring OBFL for a Module

To configure OBFL for specific modules on the switch, follow these steps

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# hw-module logging onboard module 1`	Enables all OBFL features on a module.
	`switch(config)# hw-module logging onboard module 1 cpu-hog`	Enables the OBFL CPU hog events on a module.
	`switch(config)# hw-module logging onboard module 1 environmental-history`	Enables the OBFL environmental history on a module.
	`switch(config)# hw-module logging onboard module 1 error-stats`	Enables the OBFL error statistics on a module.
	`switch(config)# hw-module logging onboard module 1 interrupt-stats`	Enables the OBFL interrupt statistics on a module.
	`switch(config)# hw-module logging onboard module 1 mem-leak`	Enables the OBFL memory leak events on a module.
	`switch(config)# hw-module logging onboard module 1 miscellaneous-error`	Enables the OBFL miscellaneous information on a module.
	`switch(config)# hw-module logging onboard module 1 obfl-log`	Enables the boot uptime, device version, and OBFL history on a module.
	`switch(config)# no hw-module logging onboard module 1`	Disables all OBFL features on a module.

Use the **show logging onboard status** command to display the configuration status of OBFL.

```
switch# show logging onboard status

    Switch OBFL Log:                                        Enabled

    Module:  6 OBFL Log:                                    Enabled
    error-stats                                             Enabled
    exception-log                                           Enabled
    miscellaneous-error                                     Enabled
    obfl-log (boot-uptime/device-version/obfl-history)      Enabled
    system-health                                           Enabled
    stack-trace                                             Enabled
```

Displaying OBFL Logs

To display OBFL information stored in CompactFlash on a module, use the following commands:

Command	Purpose
show logging onboard boot-uptime	Displays the boot and uptime information.
show logging onboard cpu-hog	Displays information for CPU hog events.
show logging onboard device-version	Displays device version information.
show logging onboard endtime	Displays OBFL logs to an end time.
show logging onboard environmental-history	Displays environmental history.
show logging onboard error-stats	Displays error statistics.
show logging onboard exception-log	Displays exception log information.
show logging onboard interrupt-stats	Displays interrupt statistics.
show logging onboard mem-leak	Displays memory leak information.
show logging onboard miscellaneous-error	Displays miscellaneous error information.
show logging onboard module *slot*	Displays OBFL information for a specific module.
show logging onboard obfl-history	Displays history information.
show logging onboard register-log	Displays register log information.
show logging onboard stack-trace	Displays kernel stack trace information.
show logging onboard starttime	Displays OBFL logs from a specified start time.
show logging onboard system-health	Displays system health information.

Default Settings

Table 59-2 lists the default system health and log settings.

Table 59-2 Default System Health and Log Settings

Parameters	Default
Kernel core generation	One module.
System health	Enabled.
Loopback frequency	5 seconds.
Failure action	Enabled.

APPENDIX A

Configuration Limits for Cisco MDS SAN-OS Release 3.x

The features supported by Cisco MDS SAN-OS have maximum configuration limits. For some of the features, we have verified configurations that support limits less that the maximum. Table A-1 lists the Cisco verified limits and maximum limits for switches running Cisco MDS SAN-OS Release 3.x.

Table A-1 Cisco MDS SAN-OS Release 3.x Configuration Limits

Feature	Verified Limit	Maximum Limit
FLOGIs or F Disc per NPV port group See "Port-Naming Conventions" section on page 4-2 for information on port groups.	114	114
NPV switches per NPV core switch	105	128
FLOGIs per line card on NPV core switch	400	400
FCNS entries per fabric	10K	10K
Device alias	8K per fabric.	20K per fabric.
Event Traps - forward via Email	1 destination.	Up to 10 destinations.
ISLB VRRP	20 per switch.	20 per switch.
VSANs	80 VSANs per physical fabric.	4000 VSANs per physical fabric.
Switches in a single MDS physical fabric or VSAN	55 switches.[1]	239 switches.
Switches in a mixed or open physical fabric or VSAN	32 switches.	239 switches.
Domains per VSAN	40 domains.	239 domains.
Zone members	16,000 zone members per physical fabric (includes all VSANs).	20,000 zone members per Physical Fabric (includes all VSANs).
Zones	8000 zones per switch (includes all VSANs).	8000 zones per switch (includes all VSANs).
Zone sets	500 zone sets per switch (includes all VSANs).	1000 zone sets per switch (includes all VSANs).
Supported hops for all major storage, server, and HBA vendors	7 hops (diameter of the SAN fabric).	12 hops.

Table A-1 Cisco MDS SAN-OS Release 3.x Configuration Limits (continued)

Feature	Verified Limit	Maximum Limit
IVR zone members	4000 IVR zone members per physical fabric.	20,000 IVR zone manbers per physical fabric in Cisco SAN-OS Release 3.0(3) and later. 10,000 IVR zone members per physical fabric prior to Cisco SAN-OS Release 3.0(3).
IVR zones	1500 IVR zones per physical fabric.	8000 IVR zones per physical fabric in Cisco SAN-OS Release 3.0(3) and later. 2000 IVR zones per physical fabric prior to Cisco SAN-OS Release 3.0(3).
IVR zone sets	32 IVR zone sets per physical fabric.	32 IVR zone sets per physical fabric.
IVR service groups	16 service groups per physical fabric.	16 service groups per physical fabric.
ISL instances per switch [2]	Up to 200 ISLs, each with 16 VSANs, for a total of 3200 port-VSAN instances. You can configure more than 200 ISLs with fewer than 16 VSANs, or fewer than 200 ISLs with more than 16 VSANs, within the total ports per VSAN instance limit of 3200.	Up to 200 ISLs, each with 16 VSANs, for a total of 3200 port-VSAN instances. You can configure more than 200 ISLs with fewer than 16 VSANs, or fewer than 200 ISLs with more than 16 VSANs, within the total ports per VSAN instance limit of 3200.
IP ports per switch	No limits.	No limits.
Fibre Channel modules vs. IPS modules per switch	No limits.	No limits.
iSCSI and iSLB sessions per IP port	500 sessions.	500 sessions.
iSCSI and iSLB sessions per switch	5000 sessions.	5000 sessions.
iSCSI and iSLB initiators supported in physical fabric	2000 initiators.	2000 initiators.
iSCSI and iSLB targets per physical fabric (virtual and initiator targets)	6000 targets.	6000 targets.

1. Certain design considerations must be met to reach this limit. We recommend that you have the large Fabric design validated by Cisco Advanced Services.
2. This is the number of trunking-enabled ISL ports multiplied by the number of VSANs in the switch.

INDEX

Symbols

Numerics

A

B

C

D

E

F

G

H

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z