



Send documentation comments to mdsfeedback-doc@cisco.com



Cisco MDS 9000 Family Configuration Guide

Cisco MDS SAN-OS Release 1.1(1a)

January, 2004

Corporate 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: 408 526-4100

Customer Order Number: DOC-7814893=
Text Part Number: 78-14893-03



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0304R)

Cisco MDS 9000 Family Configuration Guide
Copyright © 2003, Cisco Systems, Inc.
All rights reserved.



Preface **xxi**

| | |
|---|-------|
| Audience | xxi |
| Organization | xxi |
| Document Conventions | xxiii |
| Related Documentation | xxiv |
| Obtaining Documentation | xxiv |
| Cisco.com | xxiv |
| Documentation CD-ROM | xxv |
| Ordering Documentation | xxv |
| Documentation Feedback | xxv |
| Obtaining Technical Assistance | xxvi |
| Cisco.com | xxvi |
| Technical Assistance Center | xxvi |
| Cisco TAC Website | xxvii |
| Cisco TAC Escalation Center | xxvii |
| Obtaining Additional Publications and Information | xxvii |

CHAPTER 1

Product Overview **1-1**

| | |
|--|-----|
| Hardware Overview | 1-1 |
| Cisco MDS 9216 Fabric Switch | 1-1 |
| Cisco MDS 9500 Modular Directors | 1-2 |
| Software Features | 1-3 |
| High Availability | 1-3 |
| Switch Reliability | 1-3 |
| Virtual SANs | 1-4 |
| Intelligent Zoning | 1-4 |
| Trunking | 1-4 |
| PortChannels | 1-5 |
| IP Services | 1-5 |
| IP Storage | 1-6 |
| Call Home | 1-6 |
| QoS and Congestion Control | 1-7 |
| Switch Management Features | 1-7 |
| Redundant Supervisor Module Management | 1-7 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|----------------------------------|-----|
| Fabric Management | 1-8 |
| Security Management | 1-8 |
| Tools for Software Configuration | 1-9 |
| CLI | 1-9 |
| Cisco MDS 9000 Fabric Manager | 1-9 |

CHAPTER 2

Before You Begin 2-1

| | |
|--|------|
| About the Switch Prompt | 2-2 |
| About the CLI Command Modes | 2-3 |
| Understanding CLI Command Hierarchy | 2-4 |
| EXEC Mode Options | 2-5 |
| Configuration Mode | 2-6 |
| Configuration Mode Commands and Submodes | 2-6 |
| Navigating Through CLI Commands | 2-9 |
| Getting Help | 2-9 |
| Command Completion | 2-9 |
| Using the no and Default Forms of Commands | 2-10 |
| Entering CLI Commands | 2-10 |
| Viewing a Configuration | 2-10 |
| Using the File System | 2-12 |
| Setting the Current Directory | 2-12 |
| Displaying the Current Directory | 2-12 |
| Listing the Files in a Directory | 2-12 |
| Creating a New Directory | 2-13 |
| Deleting an Existing Directory | 2-13 |
| Moving Files | 2-13 |
| Copying Files | 2-14 |
| Displaying File Contents | 2-14 |
| Displaying Disk Usage | 2-14 |
| Displaying Users | 2-14 |
| Executing Commands Specified in a Script | 2-15 |
| Setting the Delay Time | 2-16 |
| Displaying the Last Line in a File | 2-16 |
| Setting the Switch's Shell Timeout | 2-16 |
| Setting the Switch's Terminal Timeout | 2-17 |
| Setting the Switch's Terminal Type | 2-17 |
| Setting the Switch's Terminal Length | 2-17 |
| Setting the Switch's Terminal Width | 2-17 |
| Displaying Terminal Settings | 2-18 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---------------------------------|------|
| Saving Command Output to a File | 2-18 |
| Sending Messages to Users | 2-18 |
| Using the ping Command | 2-18 |
| Using traceroute | 2-19 |
| Saving a Configuration | 2-19 |
| Clearing a Configuration | 2-19 |
| Role-Based CLI | 2-19 |
| Using Valid Formats and Ranges | 2-20 |

CHAPTER 3

| | |
|---|------------|
| Initial Configuration | 3-1 |
| Starting a Switch in the Cisco MDS 9000 Family | 3-2 |
| Initial Setup Routine | 3-2 |
| Preparing to Configure the Switch | 3-3 |
| Default Login | 3-3 |
| Setup Options | 3-4 |
| Assigning Setup Information | 3-5 |
| Configuring Out-of-Band Management | 3-5 |
| In-Band Management Configuration | 3-9 |
| Using the setup Command | 3-12 |
| Assigning a Switch Name | 3-12 |
| Assigning SNMP Switch Contact Information | 3-13 |
| Accessing the Switch | 3-14 |
| Where Do You Go Next? | 3-14 |
| Verifying the Module Status | 3-15 |
| Configuring Time | 3-15 |
| Configuring the Time Zone | 3-16 |
| Setting the Daylight Saving Time Adjustment | 3-16 |
| NTP Configuration | 3-18 |
| NTP Configuration Guidelines | 3-19 |
| Configuring the Management Port | 3-20 |
| Configuring Default Gateways | 3-21 |
| Disabling a Telnet Server | 3-22 |
| About Flash Devices | 3-22 |
| Formatting Flash Disks and File Systems | 3-23 |
| Initializing bootflash: | 3-23 |
| Formatting Slot0: | 3-23 |
| Working with Configuration Files | 3-24 |
| Guidelines for Creating and Using Configuration Files | 3-24 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|------|
| Viewing Configuration Files | 3-24 |
| Downloading Configuration Files to the Switch | 3-25 |
| From a Remote Server | 3-25 |
| From an External Flash (slot0:) | 3-26 |
| To a Remote Server | 3-26 |
| To an External CompactFlash Disk | 3-27 |
| Saving the Configuration | 3-27 |
| Copying Files | 3-28 |
| Backing up the Current Configuration | 3-29 |
| Rolling Back to a Previous Configuration | 3-29 |
| Restoring the Configured Redundancy Mode | 3-30 |
| Deleting Files | 3-30 |
| Configuring Line Console Settings | 3-31 |
| Console Port Speed | 3-31 |
| Device Control Parameters | 3-31 |
| Configuring the COM 1 Port | 3-32 |
| Configuring CDP | 3-33 |
| Clearing CDP Configurations | 3-34 |
| Displaying CDP Protocol Settings | 3-35 |

CHAPTER 4

Configuring High Availability 4-1

| | |
|----------------------------------|-----|
| About High Availability | 4-2 |
| Switchover Mechanisms | 4-3 |
| HA Switchover | 4-3 |
| Warm Switchover | 4-3 |
| Configuring System Switchover | 4-4 |
| Switchover Guidelines | 4-4 |
| Process Restartability | 4-5 |
| Synchronizing Supervisor Modules | 4-5 |
| Displaying HA Information | 4-6 |
| Default Settings | 4-8 |

CHAPTER 5

Software Images 5-1

| | |
|---|-----|
| About Software Images | 5-2 |
| Essential Upgrade Prerequisites | 5-2 |
| Software Upgrade Mechanisms | 5-4 |
| Performing an Automated, One-Step Upgrade | 5-5 |
| The install all command | 5-5 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|------|
| Benefits of Using the install all Command | 5-5 |
| Recognizing Failure Cases | 5-6 |
| Using the install all Command | 5-7 |
| Sample install all Commands | 5-8 |
| Performing a Manual Upgrade on a Dual Supervisor Switch | 5-12 |
| Preparing for a Manual Installation | 5-12 |
| Upgrading a Loader | 5-15 |
| Upgrading the BIOS | 5-16 |
| Specifying Kickstart and System Boot Variables | 5-18 |
| Performing a System Switchover | 5-19 |
| Performing a Manual Upgrade on a Single Supervisor Switch | 5-21 |
| Quick, One-Step Upgrade | 5-24 |
| Recovering a Corrupted Bootflash | 5-25 |
| Recovery Using BIOS Setup | 5-27 |
| Recovery from the loader> Prompt | 5-30 |
| Recovery from the switch(boot)# Prompt | 5-31 |
| Recognizing Error States | 5-32 |
| Recovery for Switches with Dual Supervisor Modules | 5-33 |
| Replacing a Faulty Supervisor Module | 5-34 |
| Default Factory Settings | 5-34 |

CHAPTER 6

| | |
|------------------------------------|------|
| Managing Modules | 6-1 |
| About Modules | 6-1 |
| Supervisor Modules | 6-2 |
| Switching Modules | 6-2 |
| Verifying the Status of a Module | 6-2 |
| Viewing the State of a Module | 6-3 |
| Connecting to a Module | 6-4 |
| Reloading Modules | 6-5 |
| Reloading the Switch | 6-5 |
| Power Cycling Modules | 6-5 |
| Reloading Switching Modules | 6-5 |
| Preserving Module Configuration | 6-6 |
| Purging Module Configuration | 6-7 |
| Powering Off Switching Modules | 6-7 |
| Identifying Module LEDs | 6-8 |
| Default Supervisor Module Settings | 6-10 |

Send documentation comments to mdsfeedback-doc@cisco.com

CHAPTER 7

Managing System Hardware 7-1

- Displaying Switch Hardware Inventory 7-2
- Displaying Power Usage Information 7-5
- Configuring Power Supplies 7-6
 - Power Supply Guidelines 7-6
- Displaying Module Temperature 7-9
- Monitoring Fan Modules 7-10
- Monitoring Clock Modules 7-10
- Displaying Environment Information 7-11

CHAPTER 8

Configuring and Managing VSANs 8-1

- How VSANs Work 8-2
- VSANs Versus Zones 8-4
- Default and Isolated VSANs 8-6
 - Default VSANs 8-6
 - Isolated VSANs 8-6
- VSAN Membership 8-6
- VSAN Attributes 8-7
 - Operational State of a VSAN 8-7
- Creating and Configuring VSANs 8-7
- Assigning VSAN Membership 8-8
- Deleting VSANs 8-9
- Viewing VSAN Configurations 8-10
- Default Settings 8-11

CHAPTER 9

Configuring Interfaces 9-1

- Configuring Fibre Channel Interfaces 9-2
 - About Interface Modes 9-2
 - E Port 9-3
 - F Port 9-3
 - FL Port 9-3
 - TL Port 9-3
 - TE Port 9-4
 - SD Port 9-4
 - Fx Port 9-4
 - B Port 9-4
 - Auto Mode 9-5

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|--------------------------------------|------|
| About Interface States | 9-5 |
| Administrative States | 9-5 |
| Operational States | 9-5 |
| Reason Codes | 9-5 |
| Configuring Fibre Channel Interfaces | 9-8 |
| Configuring a Range of Interfaces | 9-8 |
| Disabling Interfaces | 9-8 |
| Configuring Interface Modes | 9-9 |
| Configuring Administrative Speeds | 9-9 |
| Configuring Interface Descriptions | 9-10 |
| Configuring Buffer-to-Buffer Credits | 9-10 |
| Configuring Receive Data Field Size | 9-11 |
| Configuring the Beacon Mode | 9-11 |
| Identifying the Beacon LEDs | 9-12 |
| Configuring Switchport Defaults | 9-13 |
| Default Settings | 9-13 |
| Configuring the Management Interface | 9-14 |
| Configuring VSAN Interfaces | 9-15 |
| Displaying Interface Information | 9-15 |
| Displaying TL Port Information | 9-22 |
| TL Port Translation Guidelines | 9-24 |

CHAPTER 10

Configuring Trunking 10-1

| | |
|-------------------------------------|------|
| About Trunking | 10-1 |
| About Trunking Protocol | 10-2 |
| Configuring Trunk Modes | 10-3 |
| Configuring Trunk-Allowed VSAN List | 10-4 |
| Trunking Configuration Guidelines | 10-6 |
| Displaying Trunking Information | 10-7 |
| Default Settings | 10-8 |

CHAPTER 11

Configuring PortChannels 11-1

| | |
|-----------------------------------|------|
| PortChannel Examples | 11-2 |
| About PortChanneling and Trunking | 11-3 |
| About Load Balancing | 11-4 |
| Creating a PortChannel | 11-5 |
| Deleting a PortChannel | 11-6 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|-------|
| Adding Interfaces to a PortChannel | 11-6 |
| Forcing an Interface Addition | 11-7 |
| Compatibility Check | 11-7 |
| Suspended State | 11-7 |
| Deleting Interfaces from a PortChannel | 11-8 |
| Considerations for PortChannel Configurations | 11-8 |
| Error Detection | 11-8 |
| Viewing PortChannel Information | 11-9 |
| Default Settings | 11-11 |

CHAPTER 12

| | |
|---|-------------|
| Configuring and Managing Zones | 12-1 |
| Zoning Features | 12-2 |
| Zoning Example | 12-3 |
| Configuring a Zone | 12-4 |
| Configuring Aliases | 12-4 |
| Zone Enforcement | 12-5 |
| Zone Sets | 12-5 |
| Active and Full Zone Set Considerations | 12-6 |
| A Default Zone | 12-9 |
| Recovering from Link Isolation | 12-10 |
| Distributing Zone Sets | 12-11 |
| Copying Zone Sets | 12-11 |
| Clearing Zone Sets | 12-11 |
| Viewing Zone Information | 12-12 |
| Default Settings | 12-15 |
| Zone Implementation | 12-16 |

CHAPTER 13

| | |
|--|-------------|
| Managing FLOGI, Name Server, and RSCN Databases | 13-1 |
| Displaying FLOGI Details | 13-1 |
| Configuring the Name Server Proxy Feature | 13-3 |
| Registering Name Server Proxies | 13-3 |
| Displaying Name Server Database Entries | 13-3 |
| Displaying RSCN Information | 13-6 |
| Sending RSCNs | 13-7 |
| Clearing RSCN Statistics | 13-7 |

Send documentation comments to mdsfeedback-doc@cisco.com

CHAPTER 14

Configuring System Security and AAA Services 14-1

| | |
|--|-------|
| Management Security Features | 14-2 |
| User Authentication | 14-2 |
| Local Authentication | 14-2 |
| RADIUS Authentication | 14-3 |
| Role-Based Authorization | 14-3 |
| Accounting | 14-3 |
| Authentication and Authorization Process | 14-4 |
| Configuring CLI Authentication Methods | 14-5 |
| Setting AAA Authentication | 14-5 |
| Enabling or Disabling Telnet Access | 14-5 |
| Displaying CLI Authentication Commands | 14-6 |
| Configuring Role-Based CLI Authorization | 14-7 |
| Configuring Rules and Features for Each Role | 14-7 |
| Displaying Role-Based CLI Information | 14-8 |
| Configuring CLI User Profiles | 14-9 |
| Creating or Updating Users | 14-9 |
| Displaying User Profile Information | 14-10 |
| Configuring CLI Accounting Parameters | 14-11 |
| Setting the Accounting Log Size | 14-11 |
| Enabling RADIUS Accounting | 14-11 |
| Displaying Accounting Configuration | 14-12 |
| Recovering Administrator Password | 14-13 |
| Configuring RADIUS Authentication | 14-14 |
| Setting the RADIUS Server Address | 14-14 |
| Setting the RADIUS Preshared Key | 14-15 |
| Setting the RADIUS Server Time-Out Interval | 14-15 |
| Setting Iterations of the RADIUS Server | 14-16 |
| Defining Vendor-Specific Attributes | 14-16 |
| VSA Format | 14-17 |
| Authorization Process | 14-17 |
| Displaying RADIUS Server Details | 14-18 |
| Configuring SSH Services | 14-18 |
| Enabling SSH Service | 14-18 |
| Generating an SSH Host Key Pair | 14-19 |
| Using the force Option | 14-19 |
| Displaying SSH Protocol Status | 14-20 |
| SNMP Security | 14-20 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|--|-------|
| SNMP Version 1 and Version 2c | 14-21 |
| SNMP Version 3 | 14-21 |
| Group-Based SNMP Access | 14-21 |
| Restricting Switch Access | 14-22 |
| Creating SNMP Groups | 14-22 |
| Creating and Modifying Users | 14-23 |
| Forcing Identical SNMP and CLI Passwords | 14-24 |
| Assigning Users to Groups | 14-24 |
| Adding or Deleting Communities | 14-24 |
| Displaying SNMP Security Information | 14-25 |
| Displaying SNMP Counter Information | 14-25 |
| Default Security Settings | 14-26 |

CHAPTER 15

Configuring Fibre Channel Routing Services and Protocols 15-1

| | |
|---|-------|
| FSPF Features | 15-2 |
| FSPF Examples | 15-2 |
| Fault Tolerant Fabric | 15-2 |
| Redundant Links | 15-3 |
| Configuring FSPF Globally | 15-3 |
| Deleting the Entire FSPF Configuration | 15-4 |
| Disabling FSPF Routing Protocols | 15-4 |
| Link State Record Defaults | 15-4 |
| Configuring FSPF for a Specific Interface | 15-5 |
| Computing Route Cost | 15-5 |
| Specifying Hello Time Intervals | 15-5 |
| Specifying Dead Intervals | 15-6 |
| Disabling FSPF for Specific Interfaces | 15-6 |
| Retransmitting Intervals | 15-7 |
| Configuring Fibre Channel Routes | 15-7 |
| Clearing FSPF Counters | 15-8 |
| Broadcast Routing | 15-9 |
| In-Order Delivery | 15-9 |
| Reordering Network Frames | 15-9 |
| Reordering PortChannel Frames | 15-10 |
| Enabling In-Order Delivery | 15-10 |
| Configuring the Drop Latency Time | 15-11 |
| Displaying Latency Information | 15-11 |
| Configuring Flow Statistics | 15-12 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|-------|
| Clearing FIB Statistics | 15-13 |
| Displaying Flow Statistics | 15-13 |
| Displaying Routing and Forwarding Information | 15-14 |
| Displaying Global FSPF Information | 15-16 |
| Displaying the FSPF Database | 15-16 |
| Displaying FSPF Interfaces | 15-18 |
| Default Settings | 15-18 |

CHAPTER 16

Configuring IP Services 16-1

| | |
|--|-------|
| Traffic Management Services | 16-2 |
| Configuring the Ethernet Management Port | 16-2 |
| Configuring the Default Gateway | 16-3 |
| Configuring the Default Network | 16-4 |
| Configuring IPFC | 16-5 |
| Configuring an IP Address in a VSAN | 16-5 |
| Disabling IP Routing | 16-5 |
| Configuring IP Static Routes | 16-6 |
| Viewing and Clearing ARPs | 16-6 |
| Displaying IP Interface Information | 16-7 |
| Configuring Overlay VSANs | 16-8 |
| Configuring Multiple VSANs | 16-10 |
| Configuring VRRP | 16-12 |
| VRRP Features | 16-12 |
| VRRP Functionality | 16-13 |
| Creating or Removing a Virtual Router | 16-14 |
| Enabling a Virtual Router | 16-14 |
| Adding an IP Address for a Virtual Router | 16-14 |
| Setting Priority for the Virtual Router | 16-15 |
| Setting the Time Interval for the Advertisement Packet | 16-15 |
| Preempting the Master Virtual Router | 16-16 |
| Configuring Authentication for the Virtual Router | 16-16 |
| Setting the Priority Based on Interface State | 16-17 |
| Displaying VRRP Information | 16-17 |
| Clearing VRRP Statistics | 16-18 |
| Configuring DNS Server | 16-19 |
| Displaying DNS Host Information | 16-20 |
| Default Settings | 16-20 |

Send documentation comments to mdsfeedback-doc@cisco.com

CHAPTER 17
Configuring IP Storage 17-1

- IP Storage Services Module 17-2
 - Verifying the Module Status 17-3
- Configuring Gigabit Ethernet Interfaces 17-4
 - About Gigabit Ethernet Interfaces 17-4
 - Basic Gigabit Ethernet Configuration 17-4
 - About VLANs for Gigabit Ethernet 17-5
 - VLAN Configuration 17-6
 - Interface Subnet Requirements 17-6
 - Managing IP Routing 17-7
 - Displaying the IP Route Table 17-7
 - Verifying Gigabit Ethernet Connectivity 17-7
 - Managing ARP Caches 17-8
 - Displaying Statistics 17-8
 - Displaying Gigabit Ethernet Interface Statistics 17-8
 - Displaying Ethernet MAC Statistics 17-9
 - Displaying DMA-Bridge Statistics 17-9
 - Displaying TCP/IP Statistics 17-10
 - Gigabit Ethernet High Availability 17-13
 - Configuring VRRP 17-13
 - Configuring Ethernet PortChannels 17-14
 - Configuring CDP 17-16
 - IPS Core Dumps 17-16
- Configuring FCIP 17-17
 - About FCIP 17-17
 - FCIP and VE Ports 17-18
 - FCIP Link 17-18
 - FCIP Profiles 17-19
 - FCIP Interface 17-19
 - Basic FCIP Configuration 17-20
 - Creating FCIP Profiles 17-20
 - Creating FCIP Links 17-21
 - Advanced FCIP Profile Configuration 17-22
 - Configuring TCP Listener Ports 17-22
 - Configuring TCP Parameters 17-22
 - Advanced FCIP Interface Configuration 17-26
 - Configuring Peers 17-27
 - Configuring Active Connection 17-28
 - Configuring the Number of TCP Connections 17-29

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|-------|
| Enabling Time Stamps | 17-29 |
| B Port Interoperability Mode | 17-30 |
| E Port Configurations | 17-32 |
| Displaying FCIP Information | 17-33 |
| FCIP High Availability | 17-35 |
| Fibre Channel PortChannels | 17-35 |
| FSPF | 17-36 |
| VRRP | 17-36 |
| Ethernet PortChannels | 17-37 |
| Ethernet PortChannels and Fibre Channel PortChannels | 17-37 |
| Configuring iSCSI | 17-38 |
| About iSCSI | 17-38 |
| Routing iSCSI Requests and Responses | 17-40 |
| Presenting Fibre Channel Targets as iSCSI Targets | 17-41 |
| Dynamic Importing | 17-41 |
| Static Importing | 17-42 |
| iSCSI Virtual Target Configuration Examples | 17-44 |
| Presenting iSCSI Hosts as Virtual Fibre Channel Hosts | 17-46 |
| Dynamic Mapping | 17-46 |
| Static Mapping | 17-47 |
| Assigning VSAN Membership to iSCSI Hosts | 17-48 |
| Access Control in iSCSI | 17-49 |
| Fibre Channel Zoning-Based Access Control | 17-49 |
| iSCSI-Based Access Control | 17-49 |
| Enforcing Access Control | 17-50 |
| User Authentication Using iSCSI | 17-51 |
| Authentication Mechanism | 17-51 |
| Displaying iSCSI Information | 17-53 |
| Displaying iSCSI Interfaces | 17-53 |
| Displaying Global iSCSI Information | 17-55 |
| Displaying iSCSI Sessions | 17-55 |
| Displaying iSCSI Initiators | 17-56 |
| Displaying iSCSI Virtual Targets | 17-59 |
| Displaying IPS Statistics | 17-60 |
| Displaying iSCSI User Information | 17-61 |
| iSCSI High Availability | 17-62 |
| Multiple IPS Ports Connected to the Same IP Network | 17-62 |
| VRRP-Based High Availability | 17-63 |
| Ethernet PortChannel-Based High Availability | 17-64 |
| iSCSI Authentication Setup Guidelines | 17-64 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|-----------------------------------|-------|
| No Authentication | 17-64 |
| CHAP with Local Password Database | 17-65 |
| CHAP with External RADIUS Server | 17-65 |
| Scenario 1 | 17-66 |
| Scenario 2 | 17-71 |
| Default IP Storage Settings | 17-76 |

CHAPTER 18

| | |
|------------------------------------|-------------|
| Configuring Call Home | 18-1 |
| Call Home Features | 18-2 |
| Call Home Configuration Process | 18-2 |
| Cisco AutoNotify | 18-3 |
| Configuring the Call Home Function | 18-3 |
| Assigning Contact Information | 18-4 |
| Configuring Destination Profiles | 18-5 |
| Configuring E-Mail Options | 18-6 |
| Configuring General E-Mail Option | 18-6 |
| Configuring SMTP Server and Ports | 18-7 |
| Enabling or Disabling Call Home | 18-7 |
| Testing Call Home Communication | 18-8 |
| Displaying Call Home Information | 18-8 |
| Default Settings | 18-9 |
| Event Triggers | 18-10 |
| Call Home Message Severity Levels | 18-11 |
| Message Contents | 18-12 |

CHAPTER 19

| | |
|--------------------------------------|-------------|
| Configuring Domain Parameters | 19-1 |
| About fcdomain Phases | 19-2 |
| Restarting the Domain | 19-3 |
| Configuring the Domain | 19-4 |
| Setting Switch Priority | 19-6 |
| Merging Stable Fabrics | 19-6 |
| Assigning Contiguous Domains | 19-7 |
| Disabling the fcdomain Feature | 19-7 |
| Setting the Fabric Name | 19-8 |
| Stopping Incoming RCFs | 19-8 |
| Enabling Persistent FC IDs | 19-9 |

Send documentation comments to mdsfeedback-doc@cisco.com

[Configuring Persistent FC IDs Manually](#) 19-10

[Purging Persistent FC IDs](#) 19-11

[Displaying fcdomain Information](#) 19-12

[Default Settings](#) 19-14

CHAPTER 20

Configuring Traffic Management 20-1

[FCC](#) 20-2

[FCC Process](#) 20-2

[Enabling FCC](#) 20-3

[Assigning FCC Priority](#) 20-3

[Displaying FCC](#) 20-3

[QoS](#) 20-4

[Enabling or Disabling Control Traffic](#) 20-4

[Displaying QoS Information](#) 20-4

[Default FCC and QoS Settings](#) 20-4

CHAPTER 21

Configuring System Message Logging 21-1

[About System Message Logging](#) 21-2

[System Log Message Format](#) 21-4

[Configuring System Message Logging](#) 21-5

[Enabling Message Logging](#) 21-5

[Configuring Console Severity Level](#) 21-5

[Configuring Module Logging](#) 21-6

[Configuring Facility Severity Level](#) 21-6

[Configuring Log Files](#) 21-6

[Configuring Syslog Servers](#) 21-7

[Outgoing Syslog Server Logging Facilities](#) 21-8

[Displaying System Message Logging Information](#) 21-8

[Default Settings](#) 21-13

CHAPTER 22

Discovering SCSI Targets 22-1

[About SCSI LUN Discovery](#) 22-1

[Starting SCSI LUN Discovery](#) 22-2

[Initiating Customized Discovery](#) 22-2

[Displaying SCSI LUN Information](#) 22-2

Send documentation comments to mdsfeedback-doc@cisco.com

CHAPTER 23

Monitoring Network Traffic Using SPAN 23-1

- About SPAN 23-2
- SPAN Sources 23-2
 - Allowed Source Interface Types 23-3
 - VSAN as a SPAN Source 23-3
 - Guidelines to Configure VSANs as a Source 23-4
- SPAN Sessions 23-5
- Specifying Filters 23-5
 - Guidelines to Specifying Filters 23-5
- SD Port Characteristics 23-5
 - Guidelines to Configure SPAN 23-6
- Configuring SPAN 23-6
 - Encapsulating Frames 23-8
 - SPAN Conversion Behavior 23-8
- Monitoring Traffic Using Fibre Channel Analyzers 23-10
 - Without SPAN 23-10
 - Using SPAN 23-10
 - Configuring Analyzers Using SPAN 23-11
 - Using a Single SD Port to Monitor Traffic 23-12
- Displaying SPAN Information 23-13
- Default Settings 23-14

CHAPTER 24

Advanced Features and Concepts 24-1

- Configuring Time Out Values 24-2
- Invoking the fctrace Feature 24-3
- Invoking the fcping Feature 24-4
- Configuring a Fabric Analyzer 24-5
 - About the Cisco Fabric Analyzer 24-5
 - Local Text-Based Capture 24-6
 - Remote Capture Daemon 24-6
 - GUI-Based Client 24-7
 - Configuring the Cisco Fabric Analyzer 24-7
 - Capturing Frames Locally 24-7
 - Sending Captures to Remote IP Addresses 24-9
 - Clearing Configured fcanalyzer Information 24-9
 - Viewing Display Filters Information 24-10
 - Display Filters 24-10
 - Defining Display Filters 24-11

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|---|-------|
| Displaying Filters Examples | 24-11 |
| Capture Filters | 24-14 |
| Permitted Capture Filters | 24-15 |
| Configuring World Wide Names | 24-16 |
| Configuring a Secondary MAC Address | 24-16 |
| Displaying WWN Information | 24-17 |
| Allocating Flat FC IDs | 24-18 |
| Enabling Loop Monitoring | 24-18 |
| Configuring the Switch for Interoperability | 24-20 |
| Configuring Interoperability | 24-21 |
| Cisco MDS 9500 Series Switches | 24-21 |
| Cisco MDS 9200 Series Switches | 24-22 |
| Verifying Interoperating Status | 24-23 |
| Cisco MDS 9500 Series Switches | 24-23 |
| Cisco MDS 9200 Series Switches | 24-26 |

CHAPTER 25

Configuring Fabric Configuration Servers 25-1

| | |
|----------------------------|------|
| About FCS | 25-2 |
| Significance of FCS | 25-3 |
| Configuring FCS | 25-3 |
| Displaying FCS Information | 25-4 |

CHAPTER 26

Monitoring System Processes and Logs 26-1

| | |
|--------------------------------|------|
| Displaying System Processes | 26-2 |
| Displaying System Status | 26-5 |
| Configuring Core and Log Files | 26-6 |
| Clearing the Core Directory | 26-7 |
| Displaying Cores Status | 26-7 |
| Configuring HA Policy | 26-8 |
| Resetting HA Statistics | 26-8 |
| Configuring Heartbeat Checks | 26-8 |
| Configuring Watchdog Checks | 26-8 |
| Configuring Upgrade Resets | 26-9 |
| Configuring Kernel Core Dumps | 26-9 |

INDEX

Send documentation comments to mdsfeedback-doc@cisco.com

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family 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 | Initial Configuration | Provides initial switch configuration options and switch access information. |
| Chapter 4 | Configuring High Availability | Provides details on the high availability feature including switchover mechanisms. |
| Chapter 5 | Software Images | Describes how to upgrade Cisco MDS 9000 Family switches, install software image files, use the Flash file system on the supervisor engine, and recover a corrupted bootflash image. |
| Chapter 6 | 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 7 | Managing System Hardware | Provides details on switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information. |

Send documentation comments to mdsfeedback-doc@cisco.com

| Chapter | Title | Description |
|------------|--|---|
| Chapter 8 | 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 9 | Configuring Interfaces | Explains port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces. |
| Chapter 10 | Configuring Trunking | Explains TE ports and trunking concepts. |
| Chapter 11 | Configuring PortChannels | Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels. |
| Chapter 12 | Configuring and Managing Zones | Defines various zoning concepts and provides details on configuring a zone set and zone management features. |
| Chapter 13 | Managing FLOGI, Name Server, and RSCN Databases | Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases. |
| Chapter 14 | Configuring System Security and AAA Services | Discusses the AAA parameters, user profiles, RADIUS authentication, SSH services, and SNMP Security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options. |
| Chapter 15 | Configuring Fibre Channel Routing Services and Protocols | Provides details and configuration information on Fibre Channel routing services and protocols. |
| Chapter 16 | Configuring IP Services | Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information. |
| Chapter 17 | Configuring IP Storage | Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together via IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol. |
| Chapter 18 | 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 19 | 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

| Chapter | Title | Description |
|------------|--|---|
| Chapter 20 | Configuring Traffic Management | Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features. |
| Chapter 21 | Configuring System Message Logging | Describes how system message logging is configured and displayed. |
| Chapter 22 | Discovering SCSI Targets | Describes how the SCSI LUN discovery feature is started and displayed. |
| Chapter 23 | Monitoring Network Traffic Using SPAN | Describes the switched port analyzer (SPAN), identifies SPAN sources, specifies filters, explains SPAN Sessions, SD port characteristics, and configuration details. |
| Chapter 24 | 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 25 | Configuring Fabric Configuration Servers | Describes how the fabric Configuration Server (FCS) feature is configured and displayed. |
| Chapter 26 | 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. |

Document Conventions

Command descriptions use these conventions:

| | |
|----------------------|---|
| boldface font | Commands and keywords are in boldface. |
| <i>italic font</i> | 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:

| | |
|-----------------------------|---|
| <code>screen font</code> | Terminal sessions and information the switch displays are in screen font. |
| boldface screen font | Information you must enter is in boldface screen font. |
| <i>italic screen font</i> | 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 following document provides related information:

Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family

Quick Start Guide for the Cisco MDS 9000 Family

Cisco MDS 9200 Series Hardware Installation Guide

Cisco MDS 9500 Series Hardware Installation Guide

Cisco MDS 9000 Family Command Reference

Cisco MDS 9000 Family Fabric Manager User Guide

Cisco MDS 9000 Family Troubleshooting Guide

Cisco MDS 9000 Family System Messages Guide

Cisco MDS 9000 Family MIB Reference Guide

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Send documentation comments to mdsfeedback-doc@cisco.com

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated regularly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual or quarterly subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order monthly or quarterly subscriptions through the online Subscription Store:

<http://www.cisco.com/go/subscription>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can e-mail your comments to mdsfeedback-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Send documentation comments to mdsfeedback-doc@cisco.com

Obtaining Technical Assistance

**Note**

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

Cisco provides Cisco.com, which includes the Cisco Technical Assistance Center (TAC) website, as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from the Cisco TAC website. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC website, including TAC tools and utilities.

Cisco.com

Cisco.com offers a suite of interactive, networked services that let you access Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

To obtain customized information and service, you can self-register on Cisco.com at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available: the Cisco TAC website and the Cisco TAC Escalation Center. The type of support that you choose depends on the priority of the problem and the conditions stated in service contracts, when applicable.

We categorize Cisco TAC inquiries according to urgency:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration. There is little or no impact to your business operations.
- Priority level 3 (P3)—Operational performance of the network is impaired, but most business operations remain functional. You and Cisco are willing to commit resources during normal business hours to restore service to satisfactory levels.
- Priority level 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively impacted by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Send documentation comments to mdsfeedback-doc@cisco.com

- Priority level 1 (P1)—An existing network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Cisco TAC Website

The Cisco TAC website provides online documents and tools to help troubleshoot and resolve technical issues with Cisco products and technologies. To access the Cisco TAC website, go to this URL:

<http://www.cisco.com/tac>

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC website. Some services on the Cisco TAC website require a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

<http://tools.cisco.com/RPF/register/register.do>

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC website, you can open a case online at this URL:

<http://www.cisco.com/tac/caseopen>

If you have Internet access, we recommend that you open P3 and P4 cases online so that you can fully describe the situation and attach any necessary files.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Before calling, please check with your network operations center to determine the Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The *Cisco Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the *Cisco Product Catalog* at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: *Internetworking Terms and Acronyms Dictionary*, *Internetworking Technology Handbook*, *Internetworking Troubleshooting Guide*, and the *Internetworking Design Guide*. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

Send documentation comments to mdsfeedback-doc@cisco.com

- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access *Packet* magazine at this URL:
<http://www.cisco.com/go/packet>
- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:
<http://www.cisco.com/go/iqmagazine>
- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:
http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html
- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:
http://www.cisco.com/en/US/learning/le31/learning_recommended_training_list.html



Product Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offer 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 advanced security and unified management features.

The Cisco MDS 9000 Family provides 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, page 1-1](#)
- [Software Features, page 1-3](#)
- [Tools for Software Configuration, page 1-9](#)

Hardware Overview

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

- Cisco MDS 9216 multilayer fabric switches contain one fixed integrated supervisor module with 16 Fibre Channel ports and an expansion slot which can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9509 multilayer directors contain two slots for supervisor modules and 7 slots for switching or services modules providing up to 224 ports (32 ports x 7 slots).
- Cisco MDS 9506 multilayer directors contain two slots for supervisor modules and 4 slots for switching or services modules providing up to 128 ports (32 ports x 4 slots).

Cisco MDS 9216 Fabric Switch

Cisco MDS 9216 fabric switches share a consistent software architecture with the Cisco MDS 9500 Series in a semi-modular chassis. They consist of the following major hardware components:

- The chassis has two slots, one of which is reserved for the supervisor module. The supervisor module provides supervisor functions and has 16 standard, Fibre Channel ports.

Send documentation comments to mdsfeedback-doc@cisco.com

- The backplane has direct plug-in connectivity to one switching module (any type).
- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to a fully loaded chassis.
- The hot-swappable fan tray has four fans managing the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and an RS-232 serial port allows switch configuration.
- Eight hot-swappable, small form-factor pluggable (SFP), LC Gigabit Ethernet ports can be configured with either short or long wavelength SFPs for connectivity up to 550m and 10km, respectively. Additionally, all ports are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

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

Cisco MDS 9500 Modular Directors

The Cisco MDS 9500 Series includes two multilayer, modular directors:

- The C Cisco MDS 9509 Director addresses the stringent requirements of large data center storage environments and consists of the following major hardware components:
 - The chassis has nine slots, two of which are reserved for the supervisor modules.
 - Up to seven hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to seven switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan tray has nine fans managing the airflow and cooling for the entire switch.
- The Cisco MDS 9506 Director addresses the stringent requirements of data center storage environments and consists of the following major hardware components:
 - The chassis has six slots, two of which are reserved for the supervisor modules.
 - Up to four hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to four switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan tray has six fans managing the airflow and cooling for the entire switch.

These modular directors have the following features:

- Two redundant, hot-swappable power supplies have AC or DC connection, each of which can supply power to the entire chassis.
- The hot-swappable fan tray has nine fans managing the airflow and cooling for the entire switch.
- Two supervisor modules ensure high availability and traffic load balancing capabilities. Each supervisor module can control the entire switch. The standby supervisor module provides redundancy in case the active supervisor module fails.

Send documentation comments to mdsfeedback-doc@cisco.com

- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and a RS-232 serial port allows switch configuration.
- Eight hot-swappable, small form-factor pluggable (SFP), LC Gigabit Ethernet ports can be configured with either short or long wavelength SFPs for connectivity up to 550m and 10km, respectively. Additionally, all ports are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

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

Software Features

This section provides an overview of the major software features of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

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. The high availability (HA) software framework includes the following:

- Provides stateful redundancy for supervisor module failure by using dual supervisor modules.
- Ensures nondisruptive software upgrade capability. See [Chapter 5, “Software Images.”](#)
- Protects against link failure using the PortChannel (port aggregation) feature. See [Chapter 11, “Configuring PortChannels.”](#) This feature is also available in Cisco MDS 9216 switches.
- Provides management redundancy using Virtual Routing Redundancy Protocol (VRRP). See the [“Configuring VRRP” section on page 16-12](#). This feature is also available in Cisco MDS 9216 switches.
- 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 Cisco MDS 9216 switches.

See [Chapter 4, “Configuring High Availability.”](#)

Switch Reliability

Switches in the Cisco MDS 9000 Family maintain internally controlled reliability services that ensure continued service with no degradation. This reliability service includes the following:

- Provides power-on self testing (POST)
- Detects errors, isolates faults, performs parity checking, and checks illegal addresses
- Enables remote diagnostics using Call Home troubleshooting features
- Displays LEDs that summarize the status of each switching module, supervisor module, power supply, and fan assembly

Send documentation comments to mdsfeedback-doc@cisco.com

Virtual SANs

VSANs (virtual SANs) enable higher security and greater scalability in Fibre Channel fabrics. VSANs provide isolation among devices that are physically connected to the same fabric. VSANs allow multiple logical SANs over a common physical infrastructure. VSANs offer the following:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN thus ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical SAN. 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 is provided by a configured backup path between the host and the switch.
- 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.

See [Chapter 8, “Configuring and Managing VSANs.”](#)

Intelligent Zoning

Zoning controls access between devices in a VSAN. Zoning accomplishes the following:

- Partitions devices that use different operating systems. In a heterogeneous environment, it is often necessary to separate servers and storage devices to avoid accidental transfer of information between devices with different operating systems. Such transfers could result in corruption or deletion of data.
- Creates logical subsets of closed user groups. Closed user groups are needed to enforce security or to separate functional areas across the fabric.
- Configures groups of devices that are separate from the rest of the fabric. Based on the assigned zone membership, devices outside the zone cannot access devices internal to the zone.
- Provides temporary access between devices (zone sets). Zone restrictions can be imposed temporarily, and then restored to revert to normal operation, if desired.

See [Chapter 12, “Configuring and Managing Zones”](#) and the [“VSANs Versus Zones”](#) section on page 8-4.

Trunking

Trunking is the term used to refer to an ISL link that carries one or more VSANs. Trunking ports receive and transmit Extended ISL (EISL) frames. EISL frames carry an EISL header containing VSAN information. Once EISL is enabled on an E port, that port becomes a TE port (see [Chapter 9, “Configuring Interfaces,”](#) and [Chapter 10, “Configuring Trunking”](#)). The trunking configuration is saved along with the interface information.

See the [“About PortChanneling and Trunking”](#) section on page 11-3.

Send documentation comments to mdsfeedback-doc@cisco.com

PortChannels

PortChannel refers to the aggregation of multiple physical Fibre Channel ports into one logical port to provide high aggregated bandwidth, load balancing, and link redundancy. Up to 16 physical ports can be aggregated into a PortChannel. PortChannels can connect to ports across switching modules. The failure of a port in one switching module does not bring down the logical PortChannel link. Specifically, a PortChannel does the following:

- Increases the aggregate bandwidth on an ISL or EISL 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 a source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) that identify the flow of the frame.
- 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 can contain up to 16 physical links and can span multiple modules for added high availability.

See [Chapter 11, “Configuring PortChannels.”](#)

IP Services

Switches in the Cisco MDS 9000 Family support the following IP services:

- IP over Ethernet —These services are limited to management traffic.
- IP over Fibre Channel (IPFC)—IPFC (RFC 2625) specifies how IP packets are transported using encapsulation schemes. By encapsulating IP frames into Fibre Channel frames, management information is exchanged among switches without requiring a separate Ethernet connection to each switch. Each switch includes:
 - Encapsulation for IP and Address Resolution Protocol (ARP) over Fibre Channel.
 - Address resolution uses the ARP server.
- IP routing services—These services include:
 - Ethernet or TCP/IP connection.
 - Static IP routing services to enable management traffic between VSANs.
 - DNS client support.
 - The Network Time Protocol (NTP) server synchronizes the system clocks of network devices.

See [Chapter 16, “Configuring IP Services.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com

IP Storage

The Cisco MDS 9000 Family IP services module integrates seamlessly into the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches. Traffic can be routed between any IP storage port and any other port on a Cisco MDS 9000 Family switch. The Cisco MDS 9000 Family IP Storage Services Module supports the full range of services available on other MDS 9000 Family Switching Modules including VSANs, security, and traffic management. It uses widely known IP to cost-effectively connect to more servers and more locations over greater distances than previously possible. It delivers both Fibre Channel over IP (FCIP) and iSCSI IP storage services and is configurable on a port-by-port basis.

- FCIP highlights
 - Simplifies data protection and business continuance strategies by enabling backup, remote replication, and disaster recovery over WAN distances using open-standard FCIP tunneling.
 - Improves utilization of WAN resources for backup and replication by tunneling up to 3 virtual Inter Switch Links (ISLs) on a single Gigabit Ethernet port.
 - Reduces SAN complexity by eliminating the need to deploy and manage a separate remote connectivity platform.
 - Preserves Cisco MDS 9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.
- iSCSI highlights
 - Extends the benefits of Fibre Channel SAN-based storage to IP-enabled servers at a lower cost point than possible using Fibre Channel interconnect alone.
 - Increases storage utilization and availability through consolidation of IP and Fibre Channel block storage.
 - Transparent operation preserves the functionality of legacy storage applications such as zoning tools.
 - Extending the Benefits of Fibre Channel SANs

See [Chapter 17, “Configuring IP Storage.”](#)

Call Home

The Call Home feature detects switch failures and sends alerts along with relevant failure information. These alerts are sent through E-mail to a user-specified customer center.

See [Chapter 18, “Configuring Call Home.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com

QoS and Congestion Control

Switches in the Cisco MDS 9000 Family provide priority queuing and flow control services.

- Priority queuing—The switches provide low and high priority quality of service (QoS) queues. While time-critical traffic is marked as high priority traffic, all other traffic is assigned to the default low priority queue.
- Fibre Channel Congestion Control (FCC)—FCC is a flow control mechanism that alleviates congestion on Fibre Channel networks. Any switch in the network can detect congestion for an output port. The switches sample frames from the congested queue and generate messages about the congestion level upstream toward the source of the congestion. The switch closest to the source, with FCC enabled, can perform one of two actions:
 - Forwards the frames as other vendor switches do.
 - Limits the flow of frames from the port causing the congestion.

See [Chapter 20, “Configuring Traffic Management.”](#)

Switch Management Features

Besides the software features already listed, there are additional management features that fall into the following categories: redundant supervisor module management, fabric management, and security management

Redundant Supervisor Module Management

Series of multilayer directors support two redundant supervisor modules. They require two supervisor modules to enforce redundant supervisor module management and high availability and restartability (see [Table 1-1](#)).

Table 1-1 Supervisor Module Options in Cisco MDS 9000 Switches

| Product | No. of Supervisor Modules | Slot | Features |
|----------------|--|---------------|--|
| Cisco MDS 9216 | One module (includes 16 Fibre Channel ports) | Slot 1 | 2-slot chassis allows one optional switching module in the other slot. |
| Cisco MDS 9506 | Two modules | Slots 5 and 6 | 6-slot chassis allows any switching module in the other four slots. |
| Cisco MDS 9509 | Two modules | Slots 5 and 6 | 9-slot chassis allows any switching module in the other seven slots. |

When the switch powers up and both supervisor modules are present, the module in slot 5 enters the active mode, while the second module in slot 6 enters the standby mode. All storage management functions occur on the active supervisor module. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

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

Send documentation comments to mdsfeedback-doc@cisco.com

Fabric Management

Switches in the Cisco MDS 9000 Family offer fabric management and control through the command-line interface (CLI) by using Telnet, SSH, or a serial console and through the Cisco MDS 9000 Fabric Manager tool by using the Simple Network Management Protocol (SNMP) services:

- SNMP versions 1, 2, and 3 are supported.
- Remote Monitoring (RMON) allows you to specify thresholds and monitor alarms on SNMP variables. Extended RMON alarms are available for supported Management Information Base (MIB) objects (see the *Cisco MDS 9000 Family MIB Reference Guide*).
- System error message logs (syslogs) are viewed through a console or Telnet session for asynchronous events such as an interface transition. Syslogs are directed to a local buffer or to an external server and are provisioned using the CLI or the Cisco Fabric Manager GUI (see the *Cisco MDS 9000 Family System Messages Guide*).

See the “CLI” section on page 1-9, the “Cisco MDS 9000 Fabric Manager” section on page 1-9, and the “SNMP Security” section on page 14-20.

Security Management

The Cisco MDS 9000 Family of switches offer strict and secure switch management options through switch access security, user authentication, and role-based access.

See [Chapter 14, “Configuring System Security and AAA Services.”](#)

Switch Access Security

Each switch can be accessed through the CLI or SNMP.

- Secure switch access—Available when you explicitly enable Secure Shell (SSH) access to the switch. SSH access provides additional controlled security by encrypting data, user IDs, and passwords. By default, Telnet access is enabled on each switch.
- SNMP access—SNMPv3 provides built-in security for secure user authentication and data encryption.

User Authentication

A strategy known as authentication, authorization, and accounting (AAA) is used to verify the identity of, grant access to, and track the actions of remote users. The Remote Access Dial-In User Service (RADIUS) protocol provides AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using the RADIUS server(s). A global, preshared, secret key authenticates communication between the RADIUS client and server. This secret key can be configured for all RADIUS servers or for only a specific RADIUS server. This kind of authentication provides a central configuration management capability.

Role-Based Access

Role-based access assigns roles or groups to users and limits access to the switch. Access is assigned based on the permission level associated with each user ID. Your administrator can provide complete access to each user or restrict access to specific read and write levels for each command.

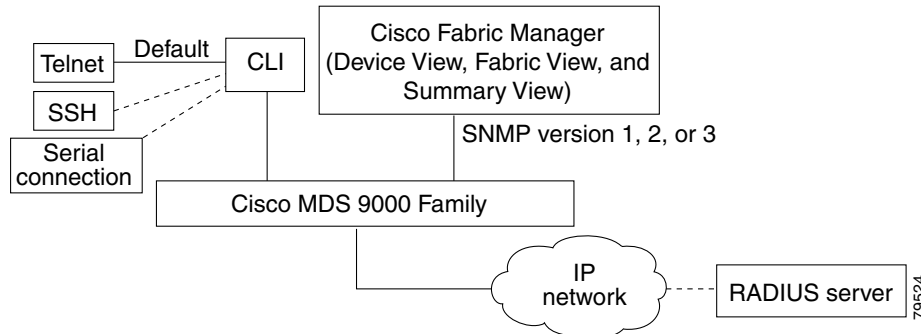
SNMP and CLI access rights are organized by roles. Each role is similar to a group. Each group of users has a specific role, and the access for that group can be enabled or disabled.

Send documentation comments to mdsfeedback-doc@cisco.com

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs: the CLI and the Cisco MDS 9000 Fabric Manager graphical user interface (GUI) from your browser (see [Figure 1-1](#)).

Figure 1-1 Tools for Configuring 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.

Cisco MDS 9000 Fabric Manager

The Cisco Fabric Manager is a Java and SNMP-based network fabric and device management tool with a GUI that displays real-time views of your network fabric and installed devices. The Cisco Fabric Manager provides three views for managing your network fabric:

- The Device View displays a continuously updated physical picture of device configuration and performance conditions for a single switch.
- The Fabric View displays a view of your network fabric, including multiple switches.
- The Summary View presents a summary view of switches, hosts, storage subsystems, and VSANs.

The Cisco Fabric Manager provides an alternative to the CLI for most switch configuration commands. The Cisco Fabric Manager is bundled with each switch in the Cisco MDS 9000 Family.

Refer to the *Cisco MDS 9000 Fabric Manager User Guide*.



Note

Resource Manager Essentials (RME) versions 3.4 and 3.5 provide support for switches in the Cisco MDS 9000 Family. Patches are available on Cisco Connection Online (<http://www.cisco.com/>).

Send documentation comments to mdsfeedback-doc@cisco.com



Before You Begin

This chapter prepares you to configure switches from the CLI. It also lists the information you need to have before you begin, and it describes the CLI command modes.

This chapter includes the following sections:

- [About the Switch Prompt, page 2-2](#)
- [About the CLI Command Modes, page 2-3](#)
- [Understanding CLI Command Hierarchy, page 2-4](#)
- [Navigating Through CLI Commands, page 2-9](#)
- [Using the File System, page 2-12](#)
- [Role-Based CLI, page 2-19](#)
- [Using Valid Formats and Ranges, page 2-21](#)

Send documentation comments to mdsfeedback-doc@cisco.com

About the Switch Prompt

If you are connected to the console port when the switch boots up, you see the output shown in [Figure 2-1](#):



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#). 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.

Figure 2-1 Output When a 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
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
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:
```


Send documentation comments to mdsfeedback-doc@cisco.com

About the 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 2-19 . | From EXEC mode, enter the config terminal command. | switch(config)# |

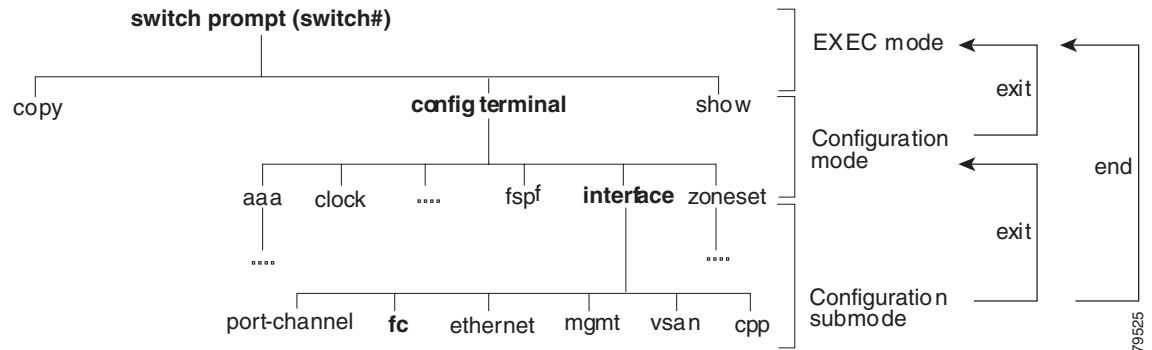
You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **config terminal** command to **conf t**.

Send documentation comments to mdsfeedback-doc@cisco.com

Understanding 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 2-2 illustrates a portion of the **config terminal** command hierarchy.

Figure 2-2 CLI Command Hierarchy Example



To execute a command, you enter the command by starting at the top level of the hierarchy. 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 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
  exit           Exit from this submode
  fcdomain       Enter the interface submode
  fspf           To configure FSPF related parameters
  no             Negate a command or set its defaults
  shutdown       Enable/disable an interface
  switchport     Configure switchport parameters
```

Send documentation comments to mdsfeedback-doc@cisco.com

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 14-3). From the 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
  clock       Manage the system clock
  config      Enter configuration mode
  copy        Copy from one file to another
  debug       Debugging functions
  delete      Remove files
  dir         Directory listing for files
  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
  install     Upgrade software
  load        Load system image
  mkdir       Create new directory
  move        Move files
  no          Disable debugging functions
  ping        Send echo messages
  purge       Deletes unused data
  pwd         View current directory
  reload      Reboot the entire box
  rmdir       Remove existing directory
  run-script  Run shell scripts
  send        Send message to all the open sessions
  setup       Run the basic SETUP command facility
  show        Show running system information
  sleep       Sleep for the specified number of seconds
  system      System management commands
  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
  undebg      Disable Debugging functions (See also debug)
  write       Write current configuration
  zone        Execute Zone Server commands
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuration Mode

Configuration mode allows you to 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 mode, zone configuration mode, and a variety of protocol-specific modes. 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

The following 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
  arp                [no] remove an entry from the ARP cache
  boot              Configure boot variables
  callhome          Enter the callhome configuration mode
  clock             Configure time-of-day clock
  end               Exit from configure mode
  exit              Exit from configure mode
  fcalias            Fcalias configuration commands
  fcanalyzer         Configure cisco fabric analyzer
  fcc               Configure FC Congestion Control
  fcdomain          Enter the fcdomain configuration mode
  fcdroplateness    Configure switch or network latency
  fcflow            Configure fcflow
  fcinterop          Interop commands.
  fcns              Name server configuration
  fcroute           Configure FC routes
  fcs               Configure Fabric Config Server
  fctimer           Configure fibre channel timers
  fspf              Configure fspf
  in-order-guarantee Set in-order delivery guarantee
  interface         Select an interface to configure
  ip                Configure IP features
  line              Configure a terminal line
  logging            Modify message logging facilities
  no                Negate a command or set its defaults
  ntp               NTP Configuration
  power             Configure power supply
  poweroff          Poweroff a module in the switch
  qos               Configure priority of FC control frames
  radius-server     Configure RADIUS related parameters
  role              Configure roles
  rscn              Config commands for RSCN
  snmp-server       Configure snmp server
  span              Enter SPAN configuration mode
  ssh               Configure SSH parameters
  switchname        Configure system's network name
  system            System config command
  telnet            Enable telnet
  trunk             Configure Switch wide trunk protocol
```

Send documentation comments to mdsfeedback-doc@cisco.com

| | |
|----------|---|
| 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 |
| 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**

When 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Table 2-2 displays the commonly used configuration submodes.

Table 2-2 Submodes Within the Configuration Mode

| Submode Name | From Configuration Mode Enter | Submode Prompt | Configured Information |
|-------------------------|--|---------------------------------------|--|
| Call Home | callhome | switch(config-callhome)# | Contact, destination, and e-mail |
| FCS Registration | fcs register | switch(config-fcs-register)# | FCS attribute registration |
| | From FCS registration submode: platform name name vsan vsan-id | switch(config-fcs-register-attr-rib)# | Platform name and VSAN ID association |
| Fibre Channel alias | fcalias name name vsan vsan-id | switch(config-fcalias)# | Alias member |
| FSPF | fspf config vsan vsan-id | switch(config-(fspf-config))# | Static SPF computation, hold time, and autonomous region |
| Interface configuration | interface type slot/port | switch(config-if)# | Channel groups, Fibre Channel domains, FSPF parameters, switch port trunk and beacon information, and IP address |
| | From the VSAN or mgmt0 (management) interface configuration submode: vrrp number | switch(config-if-vrrp)# | Virtual router (see “Creating or Removing a Virtual Router” section on page 16-14) |
| Line console | line console | switch(config-console)# | Primary terminal console |
| VTY | line vty | switch(config-line)# | Virtual terminal line |
| Role | role name | switch(config-role)# | Rule |
| SPAN | span session number | switch(config-span)# | SPAN source, destination, and suspend session information |
| VSAN database | vsan database | switch(config-vsan-db)# | VSAN database |
| Zone | zone name string vsan vsan-id | switch(config-zone)# | Zone member |
| Zone set | zoneset name name vsan vsan-id | switch(config-zoneset)# | Zone set member |

Send documentation comments to mdsfeedback-doc@cisco.com

Navigating Through CLI Commands

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.

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.

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>
fcalias          fcdomain          fcs
fcanalyzer       fcdroplacency    fcns              fctimer
fcc              fcinterop       fcroute
switch(config)# fcd<Tab>
fcdomain         fcdroplacency
switch(config)# fcd<Tab>
switch(config)# fcdomain
```

Send documentation comments to mdsfeedback-doc@cisco.com

Using 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 **zone member** command, you can undo the results:

```
switch(config)# zone name test vsan 1
switch(config-zone)# member pwnn 12:12:12:12:12:12:12:12
switch(config-zone)# no member pwnn 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 submode and return to configuration mode.

Entering CLI Commands

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 3-24).

Viewing a 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.

You can gather specific information on the entire switch configuration by issuing the relevant **show** commands. Available **show** commands for each feature are listed at the end of each chapter. Examples 2-1 to 2-3 display a few **show** command examples.

Example 2-1 Displays 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
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
0 frames input, 0 bytes, 0 discards
0 runs, 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 2-2 *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) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.

Software
  BIOS:      version 1.0.3
  loader:    version error [last 1.0(1)]
  kickstart: version 1.1(1) [build 1.1(0.94)] [gdb]
  system:    version 1.1(1) [build 1.1(0.94)] [gdb]

  BIOS compile time:      11/18/02
  kickstart image file is: bootflash:/bootimage
  kickstart compile time: 2/12/2003 11:00:00
  system image file is:   isanimage
  system compile time:    2/12/2003 12:00:00

Hardware
  RAM 1027628 kB

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

172.22.90.171 uptime is 0 days 2 hours 48 minute(s) 26 second(s)

Last reset at 669882 usecs after Thu Feb 13 07:20:41 2003
Reason: Reset Requested by CLI command reload
System version: 1.0(1)
```

Example 2-3 *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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Using the File System

The file system on a switch in the supervisor module provides a number of useful commands to help you manage software image files and configuration files.

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 directories within the internal bootflash: file system.

- The volatile: directory which provides temporary storage, and is also the default. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash (nonvolatile storage): directory which provides permanent storage. The files in bootflash are preserved through reboots and power outages.

Setting the Current Directory

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



Tip

Any file saved in the volatile: file system will be erased when the switch reboots.

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

This example changes the current directory to the mystorage directory that resides in the slot0 directory:

```
switch# cd slot0:mystorage
```

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:
```

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 file name*

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

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

Send documentation comments to mdsfeedback-doc@cisco.com

```
20971520 bytes free
20971520 bytes available
```

Creating a New 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 example deletes the directory called test at the current directory level.

```
switch# rmdir test
```

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. 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 slot0 directory 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Copying Files

The **copy** command copies a file.

This example copies the file called samplefile from the slot0 directory 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.

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 Files” section on page 3-28).

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 file_name**

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
```

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

```
switch# show file myfile
```

Displaying Disk Usage

The **show flash** command displays the disk usage of various devices.

```
switch# show flash
```

| Filesystem | 1k-blocks | Used | Available | Use% | Mounted on |
|--------------|-----------|-------|-----------|------|------------|
| none | 409600 | 37024 | 372576 | 10% | /system |
| none | 204800 | 39816 | 164984 | 20% | /var |
| none | 102400 | 2084 | 100316 | 3% | /dev/shm |
| none | 20480 | 0 | 20480 | 0% | /volatile |
| /dev/hd-cfg0 | 19976 | 1699 | 17246 | 9% | /mnt/cfg/0 |
| /dev/hd-pss | 20005 | 1473 | 17499 | 8% | /mnt/pss |

Displaying Users

The **show users** command displays all users currently accessing the switch.

```
switch# show users
```

| Username | Privilege | Time | Source |
|----------|-----------|--------------|-------------------------------|
| admin | pts/7 | Jan 12 20:56 | (10.77.202.149) |
| admin | pts/9 | Jan 12 23:29 | (modena.cisco.com) |
| admin | pts/11 | Jan 13 01:53 | (dhcp-171-71-49-49.cisco.com) |

Send documentation comments to mdsfeedback-doc@cisco.com

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 to the bootflash: directory. This section assumes that the script file resides in the bootflash: directory.

The syntax for this command is **run-script file_name**

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):
    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
```

Send documentation comments to mdsfeedback-doc@cisco.com

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 script called `lashtest-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`, the switch software executes the **discover scsi-target remote** command, and then waits for 10 seconds before executing the **show scsi-target disk** command.

Displaying the Last Line in a File

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

The syntax for this command is **tail** *<file name>* [*<number of lines>*]

```
switch# tail mylog 10
```

You see the last 10 lines of the `mylog` file.

If you specify a long file and would like to exit in the middle, enter **Ctrl-c** to exit this command.

Setting the Switch's Shell 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 from 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Setting the Switch's 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 from 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 Switch's 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 Switch's Terminal Length

To set the terminal screen length for the current session, use the **terminal length** command in EXEC mode. 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 Switch's Terminal Width

To set the terminal screen width for the current session, use the **terminal width** command in EXEC mode. 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Terminal Settings

The `show terminal` command displays the terminal settings for the current session:

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

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 3, “Initial Configuration,”](#) for information on saving and copying configuration files, and [Chapter 5, “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.

Sending Messages to Users

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

This example 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 Command

The `ping` command verifies the connectivity of a remote host or server by sending echo messages.

The syntax for this command is `ping <host or ip 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
```

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

Send documentation comments to mdsfeedback-doc@cisco.com

Using traceroute

The **traceroute** command prints the routes taken by a specified host or IP address.

The syntax for this command is **traceroute** *<host or ip address>*

```
switch# traceroute www.cisco.com
Tracing route to www.cisco.com [198.133.219.25] 30 hops max, 38 byte packets
 1  bras3-10.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
```

To abnormally terminate a traceroute session, enter **Ctrl-C**.

Saving a Configuration

To save the configuration, enter the **copy running-config startup-config** command from the EXEC mode prompt 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 Files” section on page 3-28](#).

Clearing a Configuration

To clear a startup configuration, enter the **write erase** command from the EXEC mode prompt. Once this command is issued, the switch’s startup configuration reverts to factory defaults. The running configuration is not affected. 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
```

Role-Based CLI

By default, two roles exist in all switches:

- Network operator—Has permission to view the configuration.
- Network administrator—Has permission to execute all commands and to set up to 64 permission levels based on user roles and groups (see [Chapter 14, “Configuring System Security and AAA Services”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Send documentation comments to mdsfeedback-doc@cisco.com

Using Valid Formats and Ranges



Note

Do not enter ellipsis (...), vertical bar (|), less or great (< >), bracket ([]), or braces ({ }) in command lines. These characters have special meaning in SAN-OS text strings.

Some commands require a MAC address, IP address, or IDs that must be designated in a standard format or given a range. See [Table 2-3](#).

Table 2-3 Valid Formats and Ranges

| Address | Description | Valid Format Example | Range |
|-----------------------------------|--|-------------------------|-----------------|
| MAC address | 6 bytes in hexadecimal format separated by colons (not case-sensitive) | 00:00:0c:24:d2:Fe | — |
| IP address | 32 bytes, written as 4 octets separated by periods (dotted decimal format) that are made up of a network section, an optional netmask section, and a host section. | 126.2.54.1 | — |
| VSAN | Integer that specifies the VSAN. | 7 | 1 to 4093 |
| VLAN | Integer that specifies the VLAN | 11 | 1 to 4093 |
| Port WWN (pWWN) | Eight hexadecimal numbers separated by colons (not case-sensitive). | 12:34:56:78:9A:BC:dE:F1 | — |
| Node WWN (nWWN) | Eight hexadecimal numbers separated by colons (not case-sensitive). | 12:34:56:78:9A:BC:dE:F1 | — |
| LUN | 8 bytes in hexadecimal format separated by colons. A minimum of two hex characters are acceptable. The valid format is hhhh[:hhhh[:hhhh[:hhhh]]] | 64 (100d = 64h) | — |
| FC ID | Six character hexadecimal value prepended by 0x. | 0xabc123 | — |
| Domain ID | Integer that specifies the domain. | 7 | 1 to 239 |
| Timers | Integer that specifies timers in milliseconds for latency, FC time out values (TOV). | 100 | 0 to 2147483647 |
| Switching module | Slot in which the applicable switching module resides. | 1 | 1 to 15 |
| Switch priority | Integer specifying switch priority. | 5 | 1 to 254 |
| Channel group | Integer that specifies a PortChannel group addition. | 1 | 1 to 100 |
| Fabric Shortest Path First (FSPF) | Integer that specifies the hold time (in milliseconds) before making FSPF computations. | 1000 | 0 to 65535 |
| Fabric Analyzer | The allowed range for the frame size limit in bytes. | 64 | 64 to 65536 |
| Fabric Analyzer captures | An example of 10 frames, limits the number of frames captured to 10. | 10 | 0 to 2147483647 |
| FCIP profile | Integer that specifies the FCIP profile | 101 | 1 to 255 |
| TCP retransmit time | Integer that specifies the minimum retransmit time for the TCP connection in milliseconds | 300 | 250 to 5000 |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 2-3 Valid Formats and Ranges (continued)

| Address | Description | Valid Format Example | Range |
|----------------------------|---|----------------------|-------------|
| Keepalive timeout | Integer that specifies the TCP connection's keepalive timeout in seconds. | 60 | 1 to 7200 |
| TCP retransmissions | Integer that specifies the maximum number of TCP transmissions. | 6 | 1 to 8 |
| PMTU | Integer that specifies the path MTU reset time in seconds | 90 | 60 to 3600 |
| TCP buffer size | Integer that specifies the advertised TCP buffer size in KB. | 5000 | 0 to 8192 |
| Traffic burst size | Integer that specifies the maximum burst size in KB. | 30 | 10 to 100 |
| Peer TCP port | Integer that specifies the TCP port number | 3000 | 0 to 65535 |
| Acceptable time difference | Integer that specifies the acceptable time difference in milliseconds for a packet being accepted. | 4000 | 1 to 60,000 |
| iSCSI pWWN allocation | Integer that specifies the number of pWWNs that must be allocated to an iSCSI initiator. | 2 | 1 to 64 |
| CDP refresh and hold time | Integer that specifies the refresh time interval and the hold time in seconds for the CDP protocol. | 60 | 5 to 255 |



Initial Configuration

This chapter describes how to initially configure switches so they can be accessed by other devices. This chapter includes the following sections:

- [Starting a Switch in the Cisco MDS 9000 Family, page 3-2](#)
- [Initial Setup Routine, page 3-2](#)
- [Assigning a Switch Name, page 3-12](#)
- [Accessing the Switch, page 3-14](#)
- [Where Do You Go Next?, page 3-14](#)
- [Verifying the Module Status, page 3-15](#)
- [Configuring Time, page 3-15](#)
- [Configuring the Management Port, page 3-20](#)
- [Disabling a Telnet Server, page 3-22](#)
- [About Flash Devices, page 3-22](#)
- [Formatting Flash Disks and File Systems, page 3-23](#)
- [Working with Configuration Files, page 3-24](#)
- [Copying Files, page 3-28](#)
- [Deleting Files, page 3-30](#)
- [Configuring Line Console Settings, page 3-31](#)
- [Configuring CDP, page 3-33](#)

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Step 1 Check that the switch is set for the correct AC (or DC) power voltages.
Refer to either the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for correct power voltages.

Step 2 Connect the power cord(s) to the switch.

Step 3 Connect the console port to the switch.



Note The console port is an asynchronous (async) serial port; any device connected to this port must be capable of asynchronous transmission.

Before connecting the console port, check the terminal documentation to determine the baud rate. The baud rate of the terminal must match the default baud rate (9600 baud) of the console port. Set up the terminal as follows (see the [“Configuring Line Console Settings”](#) section on page 3-31):

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

Step 4 Power on the switch. The switch boots automatically.

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 also required if you plan to configure and manage the switch.



Note The IP address must first be set up in CLI when the switch is powered up for the first time so the Cisco MDS 9000 Fabric Manager can reach the switch.

Send documentation comments to mdsfeedback-doc@cisco.com

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—You have the option to create a new login account or overwrite a preexisting account password.
- SNMPv3 user name and authentication password.
- SNMP community string.
- Switch name—This is your switch prompt.
- IP address for the switch's management interface—The management interface can be an out-of-band Ethernet interface or an in-band Fibre Channel interface.
- Subnet mask for the switch's management interface.
- The following IP addresses:
 - Destination prefix, destination prefix subnet mask, and next hop IP address, if you want to enable IP routing. Also, provide the IP address of the default network.
 - Otherwise, provide an IP address of the default gateway.
- DNS IP address (optional).
- Default domain name (optional).
- SSH service on the switch—if you wish to enable this service, then select the type of SSH key (dsa/rsa/rsa1) and number of key bits (768 to 2048).
- NTP server IP address (optional).



Note

Be sure to configure the IP routing, the IP default network address, and the IP default gateway 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.

Default Login

All Cisco MDS 9000 family switches have the network administrator as a default user (admin) and a default password (admin). You can change the default password, if required, during the initial setup process. You cannot change the default user at any time.

During the initial setup process, you have the option to configure one additional user in the network administrator role. See the [“Role-Based Authorization” section on page 14-3](#) for information of default roles and permissions.

If you change the administrator password during the initial setup process and subsequently forget this new password, you have the option to recover this password (see the [“Recovering Administrator Password” section on page 14-13](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Setup Options

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

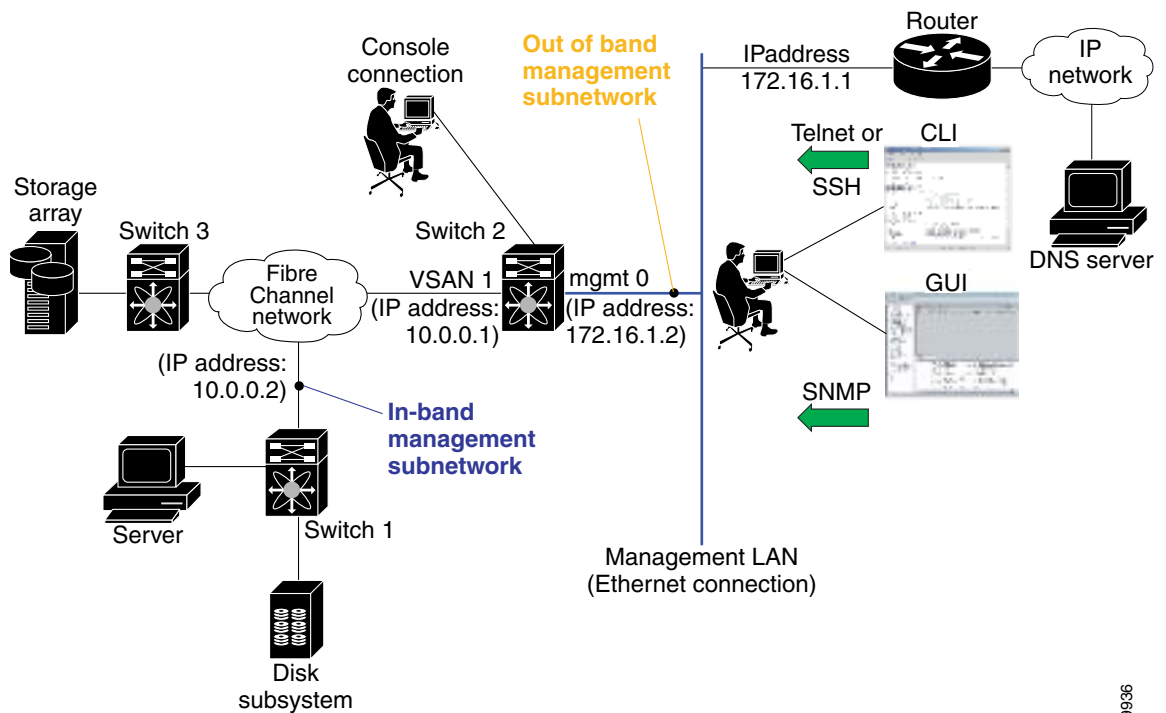


Note

Some concepts like 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 3-1](#)).
- 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 3-1](#) and [Chapter 16, “Configuring IP Services”](#)).

Figure 3-1 Management Access to Switches



79936

Send documentation comments to mdsfeedback-doc@cisco.com

Assigning Setup Information

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



Note

Type **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 is previously configured and skips to the next question.

Configuring Out-of-Band Management



Note

If you wish to configure both in-band and out-of-band configuration together, enter **Yes** in both Step 11 and Step 12 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 **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.

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. Type **Ctrl-c** at any prompt, to end the configuration process.

Step 3 Enter the new password for the administrator (admin is the default):

Enter the password for admin: **admin**

Step 4 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 14-3 for information of default roles and permissions.

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*

Send documentation comments to mdsfeedback-doc@cisco.com

Step 5 Enter **yes** (yes is the default), if you wish 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).

SNMPv3 user authentication password : *admin_pass*



Note

If this password contains less than eight characters, you must enter a new password. The same password is also used for the SNMPv3 privacy.

By default, if the admin password is at least eight characters, then the SNMP authentication password will be same as 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 6 Enter **yes** (no is the default) to configure 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 7 Enter a name for the switch.



Note

The switch name is limited to 32 alphanumeric characters.

Enter the switch name: *switch_name*

Step 8 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 IP address.

Mgmt0 IP address: *ip_address*

b. Enter the mgmt0 subnet mask.

Mgmt0 IP netmask: *subnet_mask*

Step 9 Enter **no** (no is the default) at the in-band management configuration prompt.

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

Step 10 Enter **yes** (yes is the default) to enable IP routing and default-gateway capabilities.

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

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

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

a. Enter the destination prefix.

Destination prefix: *dest_prefix*

Send documentation comments to mdsfeedback-doc@cisco.com

- b. Type the destination prefix mask.

Destination prefix mask: *dest_mask*

- c. Type the next hop ip address.

Next hop ip address: *next_hop_address*



Note

Be sure to configure the IP routing, the IP default network address, and the IP default gateway 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.

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

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

- a. Enter the default-network IP address.



Note

The default network address is the destination prefix provided in Step 10 a above.

Default network IP address: *dest_prefix*

- c. Enter **yes** (yes is the default) to configure the default-gateway (recommended).

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

- a. Enter the default-gateway IP address.

IP address of the default-gateway: *default_gateway*

- Step 11** Enter **yes** (yes is the default) to configure the DNS IP address.

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

- a. Enter the DNS IP address.

DNS IP address: *name_server*

- Step 12** Enter **yes** (default is no) to configure the default domain name.

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

- a. Enter the default domain name.

Default domain name: *domain_name*

- Step 13** Enter **yes** (yes is the default), to enable Telnet service.

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

- Step 14** Enter **yes** (no is the default) to enable the SSH service.

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

- Step 15** Enter the SSH key type (see [Generating an SSH Host Key Pair, page 14-19](#)) you would like to generate.

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

- Step 16** Enter the number of key bits within the specified range.

Send documentation comments to mdsfeedback-doc@cisco.com

Enter the number of key bits? (512 to 2048): **768**

Step 17 Enter **yes** (no is the default) to configure the NTP server.

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

a. Enter the NTP server IP address.

NTP server IP address: *ntp_server_IP_address*

Step 18 Enter **shut** (shut is the default) to configure the default switchport interface to the shut state.

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

Step 19 Enter **on** (on is the default) to configure the switchport trunk mode.

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

Step 20 Enter **permit** (deny is the default) to permit a default zone policy.

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

Permits traffic to flow to all members of the default zone.

Step 21 Review and edit the configuration that you have just entered.

Step 22 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 user admin network-admin auth md5 admin_pass priv admin_pass
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 ntp_server
system default switchport shutdown
system default switchport trunk mode on
no zone default-zone permit vsan 1-4093
```

Would you like to edit the configuration? (yes/no): **no**

Step 23 Enter **yes** (yes is default) to use and save this configuration:

Use this configuration and save it? (yes/no): **yes**



Caution

If you do not save the configuration at this point, none of your changes will be updated the next time the switch is rebooted. Ensure to type **yes** in order to save the new configuration. This will ensure that the kickstart and system boot images are also automatically configured (see [Chapter 5, “Software Images”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

In-Band Management Configuration

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 an IP address in the same subnet. A default route, pointing to the switch that provides access to the IP network should be configured on every switch in the Fibre Channel fabric (see [Chapter 8, “Configuring and Managing VSANs”](#)).



Note

If you wish to configure both in-band and out-of-band configuration together, enter **Yes** in both Step 11 and Step 12 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 **yes** to enter the setup mode.

```
---- 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.
```

```
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. Type **Ctrl-c** from any prompt, to abort the configuration process.

Step 3 Enter the new password for the administrator.

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

Step 4 Enter **no** (no is the default), if you do not wish to create other additional accounts.

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

Step 5 Enter **yes** (yes is the default), if you wish to create a SNMPv3 account.

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

a. Enter the user name.

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

By default, the SNMP user name is **admin**.

b. Enter the SNMPv3 password (minimum of 8 characters).

```
SNMPv3 user authentication password [admin_pass]: admin_pass
```

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

If this password contains less than eight characters, you must enter a new password. The same password is also used for the SNMPv3 privacy.

By default, if the admin password is at least eight characters, then the SNMP authentication password will be same as 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 6 Configure read-only or read-write SNMP community string.

- a. Enter **no** (no is the default) to avoid configuring read-only SNMP community string.

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

- b. Enter **no** (no is the default) to configure 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 7 Enter a name for the switch.

**Note**

The switch name is limited to 32 alphanumeric characters.

Enter the switch name: *switch_name*

Step 8 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 9 Enter **yes** (no is the default) at the in-band management configuration prompt.

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

- a. Enter the VSAN 1 IP address.

VSAN1 IP address: *ip_address*

- b. Enter the subnet mask.

VSAN1 IP net mask: *subnet_mask*

Step 10 Enter **yes** (yes is the default) to enable the default-gateway capabilities.

Enable ip routing capabilities? (yes/no) [y]: **yes**

- a. Enter **no** (yes is the default) to configure a static route.

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

- b. Enter **yes** (yes is the default) to configure the default network.

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

Send documentation comments to mdsfeedback-doc@cisco.com

- c. 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 11** Enter **No** (yes is the default) to configure the DNS IP address.

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

- Step 12** Enter **no** (no is the default) to skip the default domain name configuration.

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

- Step 13** Enter **no** (yes is the default), to disable Telnet service.

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

- Step 14** Enter **yes** (no is the default) to enable the SSH service.

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

- Step 15** Enter the SSH key type (see [Generating an SSH Host Key Pair, page 14-19](#)) you would like to generate.

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

- Step 16** Enter the number of key bits within the specified range.

Enter the number of key bits? (512 to 1024): **1024**

- Step 17** Enter **no** (no is the default) to configure the NTP server.

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

- Step 18** Enter **noshut** (shut is the default) to configure the default switchport interface to the up state.

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

- Step 19** Enter **auto** (on is the default) to configure the switchport trunk mode automatically.

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

- Step 20** Enter **deny** (deny is the default) to deny a default zone policy.

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

- Step 21** Review and edit the configuration that you have just entered.

- Step 22** 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 user snmp_user network-admin auth md5 snmp_pass priv snmp_pass
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
no system default switchport shutdown
system default switchport trunk mode auto
no zone default-zone deny vsan 1-4093
```

Would you like to edit the configuration? (yes/no): **no**

Send documentation comments to mdsfeedback-doc@cisco.com

Step 23 Enter **yes** (yes is default) to use and save this configuration.

Use this configuration and save it? (yes/no): **yes**



Caution

If you do not save the configuration at this point, none of your changes will be updated the next time the switch is rebooted. Ensure to type **yes** in order to save the new configuration. This will ensure that the kickstart and system boot images are also automatically configured (see [Chapter 5, “Software Images”](#)).

Using the setup Command

If you wish 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.

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.



Note

The switch name is limited to 32 alphanumeric characters.

This guide refers to a switch in the Cisco MDS 9000 Family as *switch*, and 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. |
| Step 3 | myswitch1(config)# no switchname switch(config)# | Reverts the switch name prompt to its factory default (switch#). |

Send documentation comments to mdsfeedback-doc@cisco.com

Assigning SNMP Switch Contact Information

Use the **snmp-server** command to set the contact information, switch location, and switch name. They are each limited to 32 characters (without spaces). Use the **no** form of the command to remove the system contact information. For more information on other **snmp-server** commands see the [“SNMP Security” section on page 14-20](#).

To configure contact information, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server contact NewUser switch(config)# | Assigns the contact name for the switch. |
| | switch(config)# no snmp-server contact NewUser switch(config)# | Deletes the contact name for the switch. |
| Step 3 | switch(config)# snmp-server location SanJose switch(config)# | Assigns the switch location. |
| | switch(config)# no snmp-server location SanJose switch(config)# | Deletes the switch location. |

Send documentation comments to mdsfeedback-doc@cisco.com

Accessing the Switch

After initial configuration, you can access the switch in one of three ways (see [Figure 3-2](#)):

- 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 GUI.
- Out-of-band (10/100BASE-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 GUI.

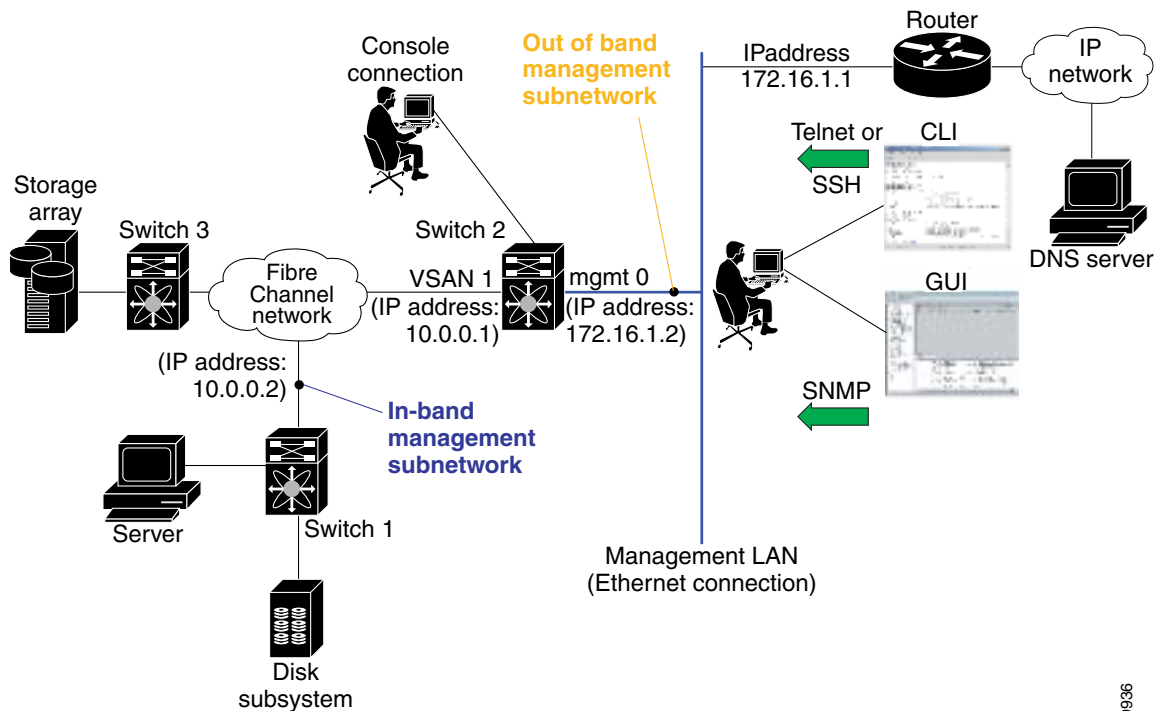


Note

To use the Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

- Serial console access—You can use a serial port connection to access the CLI.

Figure 3-2 Switch Access Options



79936

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 Element Manager and Fabric Manager GUIs.

To use the Cisco MDS 9000 Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

Send documentation comments to mdsfeedback-doc@cisco.com

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
---  -
3    16     1/2 Gbps FC Module        DS-X9016            ok
5     0     Supervisor/Fabric-1       DS-X9530-SF1-K9     active *
9    16     1/2 Gbps FC Module        DS-X9016            ok

Mod  Sw          Hw      World-Wide-Name(s) (WWN)
---  -
3    1.0(1.22)   0.0     20:81:00:05:30:00:12:5e to 20:90:00:05:30:00:12:5e
5    1.0(1.22)   0.0     --
9    1.0(1.22)   0.0     22:01:00:05:30:00:12:5e to 22:10:00:05:30:00:12:5e

Mod  MAC-Address(es)                Serial-Num
---  -
3    00-05-30-00-76-26 to 00-05-30-00-76-2a
5    00-05-30-00-53-ae to 00-05-30-00-53-b2
9    00-05-30-00-64-b6 to 00-05-30-00-64-ba
```

* this terminal session

If the status is OK or active, you can continue with your configuration (see [Chapter 6, “Managing Modules”](#)).

Configuring Time

Switches in the Cisco MDS 9000 Family use Coordinated Universal 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 12:07:50 23 September 2002
Mon Sep 23 12:07:50 UTC 2002
```

HH represents hours in military format (15 for 3 p.m.), *MM* is minutes (58), *SS* is seconds (09), *DD* is the date (02), *Month* is the month in words (August), and *YYYY* is the year (2002).



Note

The **clock** command changes are saved across system resets.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring the Time Zone

You can specify a time zone for the switch.

To specify the local time without the daylight savings 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. |

Setting the Daylight Saving Time Adjustment

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.

To enable the daylight saving time clock adjustment according to the U.S. rules, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# clock timezone <i>timezone_name hour_offset_from_UTC minute_offset_from_UTC</i> Example: switch(config)# clock timezone PST -8 0 switch(config)# no clock timezone | Offsets the time zone as specified. This example set the Pacific standard offset time as negative 8 hours and 0 minutes. Disables the timezone adjustment feature. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|---|--|
| Step 3 | <pre>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</pre> <p>Example:</p> <pre>switch(config)# clock summer-time PDT 1 Sun Apr 02:00 5 Sun Oct 02:00 60 switch(config)#</pre> | <p>Sets the daylight savings time for a specified time zone.</p> <p>The start and end values are as follows:</p> <ul style="list-style-type: none"> • week ranging from 1 through 5 • day ranging from Sunday through Saturday • month ranging from January through December <p>The daylight offset ranges from 1 through 1440 minutes which are added to the start time and deleted time from the end time.</p> <p>This example adjusts the daylight savings time for the Pacific daylight time by 60 minutes starting the first Sunday in April at 2 a.m. and ending the last Sunday in October at 2 a.m.</p> |
| | <pre>switch(config)# no clock summer-time</pre> | Disables the daylight saving time adjustment feature. |
| Step 4 | <pre>switch(config)# exit switch#</pre> | Returns to EXEC mode. |
| Step 5 | <pre>switch# show clock</pre> | Verifies the time zone configuration. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 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.

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 will be 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) act 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.

To configure NTP in a server association, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ntp server 10.10.10.10 switch(config)# | Forms a server association with a server. |
| Step 3 | switch(config)# ntp peer 10.20.10.0 switch(config)# | 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 ----- 10.20.10.2 Server 10.20.10.0 Peer | Displays the configured server and peer associations. Note A domain name will be resolved only when you have a DNS server configured. |

Send documentation comments to mdsfeedback-doc@cisco.com

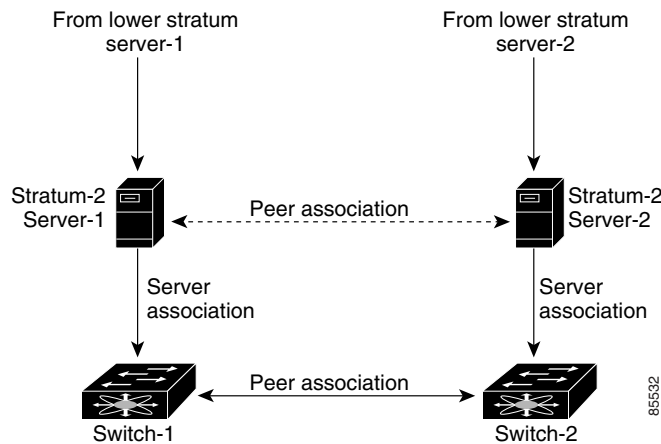
NTP Configuration Guidelines

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).
- Though a peer configured alone, will be the most accurate peer taking on the role of a server, the configured peer should be used more as a back-up support. If more than one server is present, you can have several switches point to one server, and the remaining to the another server, and then configure peer association between these two sets. This forces the clock more reliable.
- If you only have one server, it's better for all the switches have a client association with that server.

If the network is configured robustly, even a server down time will not affect well-configured switches in the network. [Figure 3-3](#) displays a network with two NTP stratum 2 servers and two switches.

Figure 3-3 NTP Peer and Server Association



In this configuration, the switches were configured as explained below:

- Stratum 2 Server 1
 - IP address -10.10.10.10
 - Stratum-2 Server-2
 - IP address -10.10.10.9
- Switch 1
 - Switch ip address -10.10.10.1
- NTP Configuration
 - NTP server 10.10.10.10
 - NTP peer 10.10.10.2
- Switch 2
 - Switch ip address -10.10.10.2
 - NTP Configuration
 - NTP server 10.10.10.9
 - NTP peer 10.10.10.1

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring the Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management port, but first you must configure some IP parameters (IP address, subnet mask) so that the switch is reachable. You can manually configure the management port interface from the CLI.



Note

Before you begin to configure the management port 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 obtain remote management access using Telnet (CLI) or SNMP (GUI), follow these steps:

| | Command | Command |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. You can also abbreviate the command to config t . The switch(config)# prompt indicates that you are in configuration mode. |
| Step 2 | switch(config)# interface type interface_string Examples: switch(config)# interface mgmt 0 | Enters the interface configuration mode on the specified interface. You can use the management Ethernet interface on the switch to configure the management interface. |
| Step 3 | switch(config)# ip address 1.1.1.0 255.255.255.0 | Enters the IP address and IP subnet mask for the interface specified in Step 2. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |
| Step 5 | switch(config-if)# exit | Returns to configuration mode. |
| Step 6 | switch(config)# ip default-gateway 1.1.1.1 | Configures the default gateway address. |

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** command.

The management port (mgmt0) is autosensing and operates as full duplex mode and 100 Mbps speed. The speed and mode cannot be configured.

When you try to shutdown 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-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-if)# shutdown force
```



Note

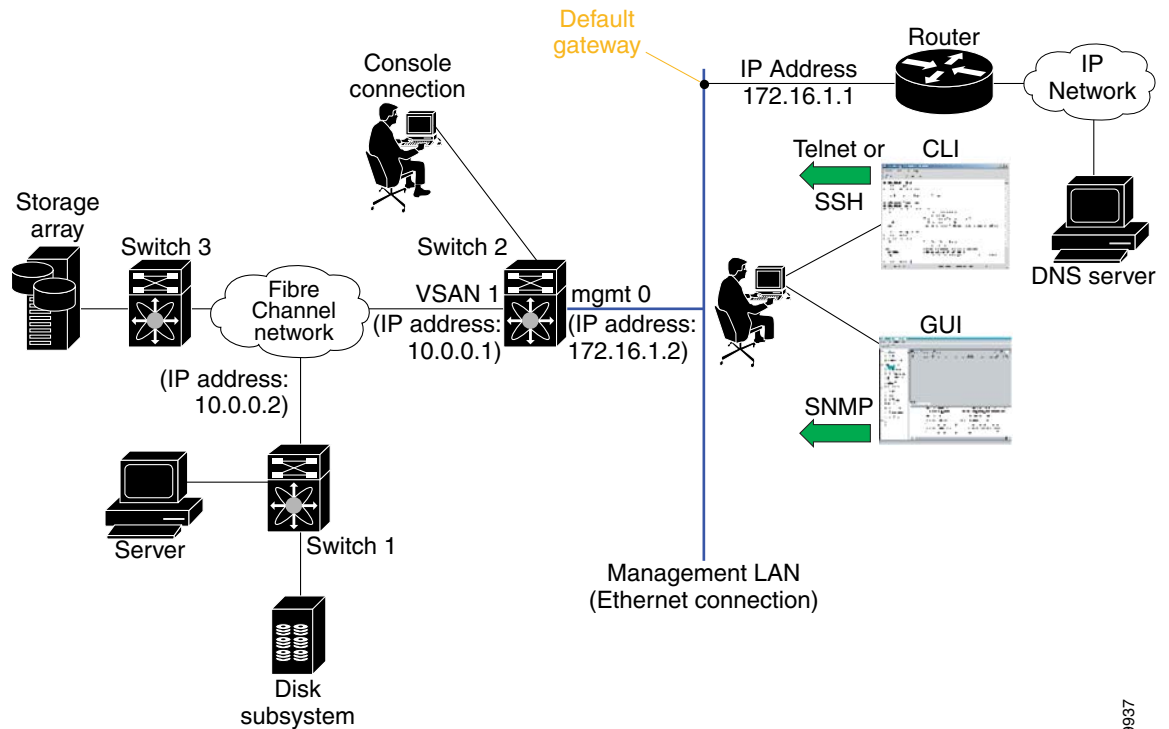
You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Default Gateways

The supervisor module sends IP packets with unresolved destination IP addresses to the default gateway (see Figure 3-4).

Figure 3-4 Default Gateway



79937

To configure the IP 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 switch(config)# | Configures the 172.16.1.1 IP address. |

Disabling a Telnet Server



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*.

| Command | Purpose |
|---|---|
| switch# config t | Enters configuration mode. |
| switch(config)# no telnet server enable updated | Disables the Telnet server. |
| switch(config)# telnet server enable updated | Enables the Telnet server if you wish to return a Telnet connection from a secure SSH connection. |

About Flash Devices

Figure 3-5 *Flash Devices in the Cisco MDS 9000 Supervisor Module*

[illegible]

Cisco MDS 9000 Family Configuration Guide

Send documentation comments to mdsfeedback-doc@cisco.com

Formatting Flash Disks and File Systems

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

Initializing 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 disk and erases all data in the bootflash: partition. The internal disk 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. If you issue an **init system** command at any time, you don't have to format the bootflash: again since 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 5, “Software Images”](#)).

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: filesystem. You can access the format bootflash: command from either the `switch#` or the `switch(boot)#` prompts.

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

Formatting Slot0:

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

You can verify if 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 the standby the `loader>` prompt or the `switch(boot)#` prompt, if the disk is inserted after booting the switch.

Send documentation comments to mdsfeedback-doc@cisco.com

Working with Configuration Files

This section describes how to work with configuration files and has the following topics:

- [Guidelines for Creating and Using Configuration Files, page 3-24](#)
- [Viewing Configuration Files, page 3-24](#)
- [Downloading Configuration Files to the Switch, page 3-25](#)
- [Saving the Configuration, page 3-27](#)
- [Copying Files, page 3-28](#)

Guidelines for Creating and Using 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.

Certain commands must be followed by a blank line in the configuration file. Without these blank lines, the commands might disconnect your Telnet session. Before disconnecting a session, the switch prompts you for confirmation. The blank line acts as a carriage return, which indicates a negative response to the prompt retaining the Telnet session.

Include a blank line after the following command in a configuration file:

```
interface mgmt0 disable
```

Viewing Configuration Files

To view the running configuration file, use the **show running-config** command:

```
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
```

To view the startup configuration file, use the **show startup-config** command:

```
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
```

Send documentation comments to mdsfeedback-doc@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 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.
- 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.

**Note**

See the “Copying Files” section on page 3-28.

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 session.
- Step 2** Configure the switch using the configuration file downloaded from the remote server using the **copy <scheme> :// <server address> system: running-config** command.
- The *scheme* is TFTP, FTP, SCP, or SFTP.
- Step 3** Specify the IP address or host name of the remote server and the name of the file to download.
- 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
```

Use the following command to download a configuration file from a remote server to the startup configuration.

```
switch# copy <scheme>://<url> nvram:startup-config
```

Send documentation comments to mdsfeedback-doc@cisco.com

From an External Flash (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 session.
 - Step 2** Locate the configuration file using the **cd** and **dir** commands. (See the “Copying Files” section on [page 3-28](#).)
 - Step 3** Configure the switch using the configuration file stored on the external CompactFlash disk using the **copy <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
```

Use the following command to download a configuration file from an external CompactFlash to the startup configuration:

```
switch copy slot0:dns-config.cfg nvram:startup-config
```

To a Remote Server

To save a configuration file to a remote server like TFTP, FTP, SCP, or SFTP, follow these steps:

-
- Step 1** Log into the switch through the console port or through a Telnet session.
 - Step 2** Save the configuration using the **copy system: running-config <scheme> :// <url>** command.
Scheme can be TFTP, FTP, SCP, or SFTP.
 - Step 3** Specify the IP address or host name of the remote server and the name of the file to download.
The configuration file is saved to 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>
```

Send documentation comments to mdsfeedback-doc@cisco.com

To an External CompactFlash Disk

To save a configuration file on an external CompactFlash disk, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Log into the switch through the console port or through a Telnet session. |
| Step 2 | Locate the configuration file using the cd and dir commands. (See the “Copying Files” section on page 3-28 .) |
| Step 3 | Save the configuration file using the copy system:running-config <source file> command. The configuration file is saved to the CompactFlash disk. |
-

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 system:startup-config slot0:dns-config.cfg
```

Saving the Configuration

After you have created a configuration, you save the configuration using the following **copy** command:

```
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.

Send documentation comments to mdsfeedback-doc@cisco.com

Copying Files

The syntax for the **copy** command follows and is explained in Table 3-1.

```
switch# copy <scheme>://<username@><server>/<file name>
<scheme>://<username@><server>/<file name>
```

Table 3-1 copy Command Syntax

| Scheme | Server | File Name |
|-------------------|---------------------------|-----------------------------------|
| bootflash | active-sup standby-sup | User-specified |
| slot0 | — | User-specified |
| volatile | — | User-specified |
| nvr | — | startup-config or snapshot-config |
| system | — | running-config |
| tftp ¹ | IP address or DNS name | User-specified |
| ftp | | |
| scp (secure copy) | | |
| sftp | | |
| core | — | 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 the active supervisor module's (sup-1 in slot 5) bootflash to the standby supervisor module's (sup-2 in slot 6) bootflash.

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



Note

Do not copy a file from an external source directly to the standby supervisor. Copy the external source to the active supervisor, and then copy the saved file to the standby supervisor.

- 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 create a running configuration copy in bootflash.

```
switch# copy system:running-config bootflash:my-config
```

- This example shows how to copy a system image file from the TFTP server to bootflash.

```
switch# copy tftp://172.16.10.100/system-237.img bootflash:system-237.img
```

- This example shows how to copy a script file from the SFTP server to volatile.

```
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.

Send documentation comments to mdsfeedback-doc@cisco.com

Backing up the Current Configuration

Before installing or migrating to any software configuration, back up the startup configuration.

- 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 backup the startup configuration copy in the bootflash: file system (ASCII file).

```
switch# copy nvram:startup-config bootflash:my-config
```

- This example shows how to backup the startup configuration to the TFTP server (ASCII file).

```
switch# copy nvram:startup-config copy tftp://172.16.10.100/my-config
```

- This example shows how to backup the running configuration to the bootflash: file system (ASCII file).

```
switch# copy system:running-config bootflash:my-config
```

Rolling Back to a Previous Configuration

All switch configurations reside in the internal bootflash: file system. If your internal bootflash: file system is corrupted, you could potentially lose your configuration. Save and back up your configuration file periodically.

- 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.

- This example shows how to roll back to a configuration copy that was previously saved in the bootflash: file system (ASCII file).

```
switch# copy bootflash:my-config nvram:startup-config
```

- This example shows how to roll back to a configuration copy that was previously saved in a TFTP server (ASCII file).

```
switch# copy tftp://172.16.10.100/my-config nvram:startup-config
```



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 to match the new system image.

Send documentation comments to mdsfeedback-doc@cisco.com

Restoring the Configured Redundancy Mode



Tip

If you have configured the **combined** mode as the redundancy mode for power supplies on a Cisco MDS 9509 switch, exert care when using the sequence of the **write erase** and **reload** commands before rolling back to a saved configuration.

As a result of issuing the **write erase** command and the **reload** command, you restore the switch settings to its factory defaults. This sequence also restores the redundancy mode setting for the power supplies back to the **redundant** mode (default).

Depending on the types of power supplies, 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 “[Configuring Power Supplies](#)” section on page 7-6).

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).

When using the sequence of the **write erase** and **reload** commands before rolling back to a saved configuration, follow these steps:

-
- Step 1** Manually change the **redundant** mode configuration to **combined** mode, if originally configured as such.
 - Step 2** Wait until all modules are back online—the module status displays **ok** in response to the **show module** command.
 - Step 3** Rollback to the saved configuration (see the “[Rolling Back to a Previous Configuration](#)” section on page 3-29).
-

Deleting Files

To delete files on a Flash device, follow these steps:

| | Command | Purpose |
|---------------|---|---------------------------------|
| Step 1 | switch# delete [device:]filename | Deletes files from a directory |
| Step 2 | switch# dir [device:][filename] | Verifies the files are deleted. |

- This example shows how to delete a file from the bootflash: directory (assuming you are already in the bootflash: 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 file named test from the Flash card inserted in slot 0.

```
switch# delete slot0:test
```

- This example deletes the entire my-dir directory:

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

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Line Console Settings

Console ports on Cisco switches are set up for quick and easy access through any standard RS-232 data terminal equipment (DTE) device.

You can perform the configuration specified in this section only if you are connected to the serial console.



Note

If you plan on connecting a modem to the console port of a switch in the Cisco MDS 9000 Family, first refer to the Console Port Issues section of the [Modem-Router Connection Guide](http://www.cisco.com) (www.cisco.com).

Console Port Speed

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.

For the purposes of this document, the default console port speed of 9600 baud is assumed.



Caution

The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To configure console port speed, 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 is 9600 baud and the range is from 110 to 115,200 baud. |

If you specify an invalid speed, you will receive the following error message:

```
switch(config-console)# speed 111
Error: 111 is not supported speed
Supported speed are 110, 150, 300,2400, 4800, 9600, 19200, 28800, 38400, 57600 (56K), and 115200
```

Device Control Parameters

Be sure to set the following values for the device control parameters when setting up the terminal:

- 8 data bits
- 1 stop bit
- No parity

You can change these parameters to meet the requirements of the terminal or host to which you are attached.

Send documentation comments to mdsfeedback-doc@cisco.com

To configure line control parameters, follow these steps:

| | Command | Command |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line console | Enters the line console configuration mode. |
| Step 3 | 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 4 | switch(config-console)# stopbits 1 | Configures the stop bits for the console connection. The default is 1 stop bits and the valid values are 1 or 2 stop bits. |
| Step 5 | 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. |

Configuring the COM 1 Port

The COM 1 port in Cisco MDS 9000 Family switches enables you to connect to an external serial communication device, such as a modem.

In many situations, you may need to configure the switch using a modem that is physically connected to the auxiliary (COM 1) port of the switch.



Note

If you plan on connecting a modem to the COM 1 port of a switch in the Cisco MDS 9000 Family, first refer to the [Modem-Router Connection Guide](http://www.cisco.com) (www.cisco.com).

To configure the COM 1 port parameters, 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 COM 1 port configuration mode. |
| Step 3 | switch(config-com1)# speed 9600 | Configures the port speed for the COM 1 connection. The default is 9600 baud and the range is from 110 to 115, 200 baud. 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. |
| 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 over CLI and SNMP.

CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. 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 protocol, 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 8/8 switch(config-if)# | Configures the Gigabit Ethernet interface for the module in slot 8 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 switch(config)# | 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 switch(config)# | 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 switch(config)# | 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 switch(config)# | Reverts the version to the factory default of v2. |

Clearing CDP Configurations

To clear CDP traffic counters for all interfaces use the **clear cdp counters** command. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces):

```
switch# clear cdp counters
switch#
```

To clear neighboring CDP entries for all interfaces, use the **clear cdp table** command. 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
switch#
```

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying CDP Protocol Settings

Use the **show cdp** command to display CDP entries. See Examples 3-1 to 3-11.

Example 3-1 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 3-2 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 3-3 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Example 3-4 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 3-5 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 3-6 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 3-7 Displays CDP Neighbors (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 | 8/11 |
| 069038747(Kiowa3 | mgmt0 | 156 | T S | WS-C5500 | 6/20 |
| 069038747(Kiowa3 | mgmt0 | 158 | T S | WS-C5500 | 5/22 |

Example 3-8 Displays CDP Neighbors (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): 8/11
Holdtime: 132 sec
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
Version:
WS-C5500 Software, Version MpsSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems
Advertisement Version: 1
```

Example 3-9 Displays the Specified CDP Neighbor (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 3-10 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 3-11 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
```

Send documentation comments to mdsfeedback-doc@cisco.com



Configuring High Availability

This chapter provides details on the high availability feature that is available on switches with two supervisor modules. It includes the following sections:

- [About High Availability, page 4-2](#)
- [Switchover Mechanisms, page 4-3](#)
- [Configuring System Switchover, page 4-4](#)
- [Switchover Guidelines, page 4-4](#)
- [Process Restartability, page 4-5](#)
- [Synchronizing Supervisor Modules, page 4-5](#)
- [Displaying HA Information, page 4-6](#)
- [Default Settings, page 4-8](#)

Send documentation comments to mdsfeedback-doc@cisco.com

About 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. The high availability (HA) software framework provides for the following:

- Ensures nondisruptive software upgrade capability. See [Chapter 5, “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 Cisco MDS 9216 switches.
- Protects against link failure using the PortChannel (port aggregation) feature. See [Chapter 11, “Configuring PortChannels.”](#)
- Provides management redundancy using Virtual Routing Redundancy Protocol (VRRP). See the [“Configuring VRRP” section on page 16-12.](#)
- Switchability—When 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 in the two center slots (sup-1 and sup-2). 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 module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Send documentation comments to mdsfeedback-doc@cisco.com

Switchover Mechanisms

When the active supervisor module fails, the standby module automatically takes over. You can also issue a **system switchover** command to manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a **system switchover** is issued (switchover process has started) another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

To determine version compatibility between switch images, use the **show system switchover impact** command.

```
switch# show system switchover impact
effects of switchover to 1.0(3)
switchover type:      HA
impact on modules:    module(s) will not reset

per module compatibility with system image 1.0(3) running on standby supervisor:
module  compatibility check
1       module running version 1.0(3): is compatible
9       module running version 1.0(3): is compatible
```

Two switchover modes are available in the Cisco MDS 9000 Family: HA switchover (default) or warm switchover.



Note

If the images are not compatible, an HA switchover is not possible.

HA Switchover

When a **show system redundancy status** or a **show module** command displays the `HA-standby` state for the standby supervisor module, an HA switchover (default) is possible. An HA switchover has the following characteristics:

- Is stateful (nondisruptive) since control traffic is not impacted
- Does not impact data traffic since the switching modules are not impacted
- Switching modules are not reset

This is the best possible scenario because there is no system downtime.

Warm Switchover

When a **show system redundancy status** or a **show module** command displays the `warm standby` state for the standby supervisor module, a warm switchover is possible. A warm switchover has the following characteristics:

- Is stateless (disruptive) since control traffic will be impacted.
- Impacts data traffic since switching modules will be impacted.
- Switching modules are reset with an significantly reduced bring up time

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring System Switchover

By default, the system uses a HA switchover. When two supervisor modules are available on the system, you can switch over from the active to the standby supervisor module using a **HA** (nondisruptive) or **warm** (disruptive) switchover. In the HA switchover mode, a HA switchover is performed where possible. If HA switchover is not possible, the warm switchover mode is attempted. If warm switchover mode is configured, then HA switchover is disabled.



Caution

Switching from HA to warm or warm to HA modes cause the standby supervisor module to reset.

To define the switchover mechanism in a switch, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# system switchover warm switch(config)# | Configures the switch to perform a stateless (disruptive) switchover the next time a switchover occurs (EXEC mode command or in response to a failure). |
| | switch(config)# system switchover HA switch(config)# | Reverts the switch settings to perform the default stateful (nondisruptive) switchover the next time a switchover occurs (EXEC mode command or in response to a failure). |
| | or issue the following command: switch(config)# no system switchover switch(config)# | Restores the default settings (HA switchover). |

Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- Use the **system switchover** command when you need to upgrade the software in a dual supervisor switch (see the “[Performing a System Switchover](#)” section on page 5-19).
- The **system switchover** command returns the following message when the standby supervisor is not present in the switch:

```
switch# system switchover
Failed to switchover: (supervisor has no standby)
```

- You can only perform a switchover when the switch has two supervisor modules functioning in the switch. Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.
- Verify 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    16     1/2 Gbps FC Module        DS-X9016             ok
5     0     Supervisor/Fabric-1       DS-X9530-SF1-K9      active *
6     0     Supervisor/Fabric-1       DS-X9530-SF1-K9      HA-standby
8    32     1/2 Gbps FC Module        DS-X9032             ok
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.0(0.253)    1.0          20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de
5    1.0(0.253)    1.0          --
6    1.0(0.253)    1.0          --
8    1.0(0.253)    1.0          20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de

```

```

Mod  MAC-Address(es)          Serial-Num
---  -
2    00-05-30-00-0f-e4 to 00-05-30-00-0f-e8  jab0636063v
5    00-05-30-00-19-66 to 00-05-30-00-19-6a  jab06370593
6    00-05-30-02-20-02 to 00-05-30-02-20-06  jab06371593
8    00-05-30-00-1a-12 to 00-05-30-00-1a-16  jab06370574

```

* 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. If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that the 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 [“Displaying System Processes” section on page 26-2](#).

Synchronizing Supervisor Modules

The **system auto-sync image** option is disabled by default on switches in the Cisco MDS 9000 Series. This command can only be operational if the following cases apply:

- the **system switchover HA** command is configured.
- two supervisor modules are up and running

You can synchronize the standby supervisor module software image with the bootflash image using the **system auto-sync image** command in configuration mode. The current running image and configuration files are synchronized from the active to the standby supervisor module (see the [“Specifying Kickstart and System Boot Variables” section on page 5-18](#)).



Note

If both supervisors modules are running the same software image, the **system auto-sync image** command will have no effect.

Send documentation comments to mdsfeedback-doc@cisco.com

To enable or disable automatic synchronization, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# system auto-sync image switch(config)# | Enables automatic synchronization. |
| | switch(config)# no system auto-sync image Automatic synchronization of BOOT and KICKSTART is now disabled. | Disables automatic synchronization (default). |

When you log into the switch after the basic upgrade, the standby supervisor module synchronizes its image automatically with the running image on the active supervisor module. To upgrade the image, you must disable this option. By disabling this option, you are ensuring that the synchronization does not take place with undesired images. Enabling this option synchronizes the running image on both supervisor modules.



Note

During a synchronization, the boot variables are not synchronized. The boot variables are independent of the two supervisor modules (see [“Performing a System Switchover”](#) section on page 5-19).

Use the **show auto-sync** command to view the status of the auto-sync configuration. See [Example 4-1](#).

Example 4-1 Displays Auto Synchronization Status

```
switch# show system auto-sync
auto-sync is disabled
auto-sync not started
```

You can view the output of the **show system redundancy** command to verify if HA switchover and automatic synchronization are enabled and operational.

Displaying HA Information

Use the **show system redundancy status** command to view the high availability status of the system. See [Example 4-2](#). Tables 4-1 to 4-3 explain the possible redundancy, supervisor, and internal states output in this command.

Example 4-2 Displays Redundancy Status

```
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
```


Send documentation comments to mdsfeedback-doc@cisco.com

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is `Active` with `HA standby` and of the other supervisor module is `HA standby`, the switch is operationally HA and can do automatic synchronization.
- If the internal state of one supervisor module is `Active` with `warm standby` and of the other supervisor module is `Warm standby`, the switch is operationally warm and cannot do automatic synchronization.
- If the internal state of one of the supervisor modules is `none` the switch cannot do automatic synchronization.

Table 4-1 lists the possible values for the redundancy states.

Table 4-1 Redundancy States

| State | Description |
|--------------|--|
| Not present | The supervisor module is not present or is not plugged into the switch. |
| Initializing | The diagnostics have passed and the configuration is being downloaded. |
| Active | This module is the active supervisor module and the switch is ready to be configured. |
| Standby | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3). |
| 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 switch is intentionally shut down for debugging purposes. |
| At BIOS | The module 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 4-2 lists the possible values for the Supervisor state.

Table 4-2 Supervisor States

| State | Description |
|--------------|--|
| Active | This module is the active supervisor module and the switch is ready to be configured. |
| HA standby | This module is the standby supervisor module and the HA switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3). |
| Warm standby | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3). |
| 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 4-3 lists the possible values for the internal state.

Table 4-3 Internal States

| State | Description |
|--------------------------------|--|
| Warm standby | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “HA Switchover” section on page 4-3). |
| HA standby | This module is the standby supervisor module and the HA switchover mechanism is enabled (see the “HA Switchover” section on page 4-3). |
| Active with no standby | This module is the active supervisor module, and the second supervisor module is not present in the switch. |
| Active with HA standby | This module is the active supervisor module and the switch is ready to be configured. The standby module is in the HA-standby state. |
| Active with warm standby | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “HA Switchover” section on page 4-3). |
| Shutting down | The switch is being shut down. |
| Warm switchover in progress | The switch is in the process of changing over to the warm switchover mechanism. |
| 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. |
| HA synchronization in progress | The standby supervisor module is in the process of synchronizing its supervisor modules. |
| Standby (failed) | The standby supervisor module is not functioning. |
| Active with failed standby | This module is 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. |

Default Settings

Table 4-4 lists the default settings for high availability features.

Table 4-4 Default High Availability Setting

| Parameters | Default |
|-----------------|---------|
| Switchover mode | HA |



Software Images

This chapter describes how to install and upgrade software images. 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 and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules. In the dual supervisor scenario, the standby supervisor module should be updated first.

This chapter includes the following sections:

- [About Software Images, page 5-2](#)
- [Essential Upgrade Prerequisites, page 5-2](#)
- [Software Upgrade Mechanisms, page 5-4](#)
- [Performing an Automated, One-Step Upgrade, page 5-5](#)
- [Performing a Manual Upgrade on a Dual Supervisor Switch, page 5-12](#)
- [Performing a Manual Upgrade on a Single Supervisor Switch, page 5-21](#)
- [Quick, One-Step Upgrade, page 5-24](#)
- [Recovering a Corrupted Bootflash, page 5-25](#)
- [Recovery for Switches with Dual Supervisor Modules, page 5-33](#)
- [Replacing a Faulty Supervisor Module, page 5-34](#)
- [Default Factory Settings, page 5-34](#)

Send documentation comments to mdsfeedback-doc@cisco.com

About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The 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 which 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 (see the [“Specifying Kickstart and System Boot Variables” section on page 5-18](#)). Both images are not always required for each install.



Note

Unless explicitly stated, the software install procedures in this section apply to any switch in the Cisco MDS 9500 Series and the Cisco MDS 9200 Series.

Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- 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 will be 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 modules or bootflash: (internal to the switch). Use the **dir** command to ensure that the required free space is available for the image files to be copied.

 - Standby supervisor module bootflash: directory (see the [“Connecting to a Module” section on page 6-4](#)).
 - 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 IP address for the 10/100 BASE-T Ethernet port connection (interface mgmt0).
 - 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. Verify connectivity to the remote server using the **ping** command.
- Images
 - The specified system and kickstart images must be 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.

Send documentation comments to mdsfeedback-doc@cisco.com

- Images can be retrieved in one of two ways:

Local—images are locally available on the switch (the **install all** command uses the specified local images).

Remote—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.

- Commands

- Verify that the **auto-sync** option is disabled.
- The switchover mode, if configured, must be configured as HA standby (see the “[Switchover Mechanisms](#)” section on page 4-3).
- 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 “[The install all command](#)” section on page 5-5).

Only one **install all** command can be running on a switch at any time.

No other command can be issued while running the **install all** command.

The **install all** command cannot be performed on the standby supervisor module—it can only be issued on the active supervisor module.

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 abort at this point.

**Note**

When you issue the **install all** command, the switch displays a summary of changes that will be made to your configuration and waits for your authorization to continue executing the command process.

Send documentation comments to mdsfeedback-doc@cisco.com

Software Upgrade Mechanisms

The Cisco MDS SAN-OS software, designed for mission-critical high availability environments, provides the ability to upgrade software without any disruptions. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9509 Director, it is highly recommended that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9500 Series and the Cisco MDS 9200 Series using one of three mechanisms:

- Automated, one-step upgrade using the **install all** command (see the [“Performing an Automated, One-Step Upgrade” section on page 5-5](#)).
- Manual, step-by-step upgrade (see the [“Performing a Manual Upgrade on a Dual Supervisor Switch” section on page 5-12](#)).
- Quick, one-step upgrade using the **reload** command. This upgrade is disruptive (see the [“Quick, One-Step Upgrade” section on page 5-24](#)).

For nondisruptive upgrades, use the automated one-step upgrade or the manual step-by-step upgrade. In some cases, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor system with kickstart image changes
- A single supervisor system with incompatible system software images
- A dual supervisor system with incompatible system software images



Tip

The **install all** command compares and presents the results of the compatibility before proceeding with the installation. You have the opportunity to exit, if you do not want to proceed with these changes.

Send documentation comments to mdsfeedback-doc@cisco.com

Performing an Automated, One-Step Upgrade

You can perform an automated upgrade on any switch in the Cisco MDS 9200 Series or the Cisco MDS 9000 Series using the **install all** command.

The install all command

The **install all** command upgrades all modules in any Cisco MDS 9000 Family switch. [Figure 5-1](#) provides an overview of the switch status before and after issuing the **install all** command.

Figure 5-1 The install all Command Effect

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.



Note

The **install all** command is only effective on switches running Cisco MDS SAN-OS Release 1.0(3) and later releases.

Benefits of Using the install all Command

The **install all** command provides the following benefits:

- You can upgrade the entire switch using just one command.
- You will receive descriptive information on the intended changes to your system before you issue the command.
- You have the option to cancel the command. Once the effects of the command are presented, you can choose to 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. The manual install procedure involves several commands and processes and takes longer.

Send documentation comments to mdsfeedback-doc@cisco.com


Tip

The automated process is described in [Performing an Automated, One-Step Upgrade, page 5-5](#). The manual process is described in [Performing a Manual Upgrade on a Dual Supervisor Switch, page 5-12](#) through [Performing a Manual Upgrade on a Single Supervisor Switch, page 5-21](#).

- You can view the progress of this command on the console screen of the supervisor module(s).
- The image integrity is automatically checked by the **install all** command. This includes the running kickstart and system images.
- A platform validity check is performed to verify that a wrong image is not used—for example, to check if an MDS 9509 image is used to upgrade an MDS 9216 switch.
- Use the **Ctrl-c** escape sequence to gracefully abort the **install all** command. The command sequence completes the update step in progress and returns to the switch prompt. (Other upgrade commands cannot be aborted using **Ctrl-c**.)
- After issuing the **install all** command, if any step in the sequence fails, the command will complete the step in progress and abort.

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 aborts. In such cases, you can verify the problem on the affected switching module and manually upgrade the other switching modules.

Recognizing Failure Cases

The following situations will cause the **install all** command to abort:

- If the **auto-sync** option is enabled.
- If the standby supervisor module bootflash: directory does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the **install all** command is issued on the standby supervisor module.
- If the fabric or switch is configured while the upgrade is in progress.
- If a module is removed or inserted 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.


Caution

Avoid aborting the switch progress after issuing the **install all** command. If the **install all** command is aborted, be sure to verify the state of the switch at every level. If you issue the command and then abort, you may need to use the manual procedure commands to complete the upgrade. This necessity depends on the state of the switch at the time of the failure.

Send documentation comments to mdsfeedback-doc@cisco.com

Using the install all Command

To perform an automated software upgrade on any switch, follow these steps:

- Step 1** Log into the switch through the console port of the active supervisor.
- Step 2** Create a backup of your existing configuration file, if required (see the [“Working with Configuration Files”](#) section on page 3-24).
- Step 3** Perform the upgrade by issuing the **install all** command on the active supervisor module.

```
switch# install all system bootflash:system-image kickstart bootflash:kickstart-image
Image verification is in progress, please wait.
This command is going to install system image m9500-sflek9-mz.1.0.3.bin
and kickstart image m9500-sflek9-kickstart-mz.1.0.3.upg.bin on this system
The command will:
- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the standby supervisor
- Perform a HA Switchover
- Perform a hitless upgrade of module 1, 2, 3, 4, 7, 8, 9
```

- Step 4** Enter **yes** if you wish to continue with the installation.

```
Do you want to continue y/n ? [n] :y

Image synchronization is in progress, please wait.

Installing Loader, please wait.
Installing Loader on module 5 ... successful
Installing Loader on module 6 ... successful

Installing BIOS, please wait.
Installing BIOS on module 1 ... not required (same version)
Installing BIOS on module 2 ... not required (same version)
Installing BIOS on module 3 ... not required (same version)
Installing BIOS on module 4 ... not required (same version)
Installing BIOS on module 5 ... not required (same version)
Installing BIOS on module 6 ... not required (same version)
Installing BIOS on module 7 ... not required (same version)
Installing BIOS on module 8 ... not required (same version)
Installing BIOS on module 9 ... not required (same version)

Updating boot variables .. successful
Saving configuration, please wait.
Reload of the standby supervisor is in progress, please wait
Success, the standby supervisor is online and ready to takeover
```

- Step 5** Exit this terminal session and open a new session to view the upgraded supervisor module using the **show module** command.



Note

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.

Send documentation comments to mdsfeedback-doc@cisco.com

Sample install all Commands

[Example 5-1](#) displays the result of the **install all** command if the system and kickstart files are specified in the command line.

Example 5-1 Files Specified Locally

```
switch# install all system bootflash:system_image kickstart bootflash:kickstart_image
```

```
Image verification is in progress, please wait.
This command is going to install system image m9500-sflek9-mz.1.0.3.bin
and kickstart image m9500-sflek9-kickstart-mz.1.0.3.upg.bin on this system
The command will:
- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the standby supervisor
- Perform a HA Switchover
- Perform a hitless upgrade of module 1, 2, 3, 4, 7, 8, 9
```

```
Do you want to continue y/n ? [n] : y
```

```
Image synchronization is in progress, please wait.
```

```
Installing Loader, please wait.
Installing Loader on module 5 ... successful
Installing Loader on module 6 ... successful
```

```
Installing BIOS, please wait.
Installing BIOS on module 1 ... not required (same version)
Installing BIOS on module 2 ... not required (same version)
Installing BIOS on module 3 ... not required (same version)
Installing BIOS on module 4 ... not required (same version)
Installing BIOS on module 5 ... not required (same version)
Installing BIOS on module 6 ... not required (same version)
Installing BIOS on module 7 ... not required (same version)
Installing BIOS on module 8 ... not required (same version)
Installing BIOS on module 9 ... not required (same version)
```

```
Updating boot variables .. successful
Saving configuration, please wait.
Reload of the standby supervisor is in progress, please wait
Success, the standby supervisor is online and ready to takeover
```

[Example 5-2](#) 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.

Send documentation comments to mdsfeedback-doc@cisco.com



Caution

Specify the exact, complete, path of the remote location. The system will not allow you to proceed if the entire path is not accurately specified. Some examples of incomplete **install all** commands are provided below, [Example 5-2](#) provides an accurate, complete example.

```
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 5-2 Files Specified in a Remote Location

```
switch# install all
system scp://user@171.71.00.000/home/user/golden-sanity/system_image
kickstart scp://user@171.71.00.000/home/user/golden-sanity/kickstart_image
Copying
scp://user@171.71.00.000/home/user/golden-sanity/system_image to bootflash:/system_image
..
Copying
scp://user@171.71.00.000/home/user/golden-sanity/kickstart_image to
bootflash:/kickstart_image
user@171.71.00.000's password:
```

```
system_image 100% |*****| 19941 KB
```

Image verification is in progress, please wait.
This command is going to install system image m9500-sflek9-mz.1.0.3.bin
and kickstart image m9500-sflek9-kickstart-mz.1.0.3.upg.bin on this system

The command will:

- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the standby supervisor
- Perform a HA Switchover
- Perform a hitless upgrade of module 1, 2, 3, 4, 7, 8, 9

Do you want to continue y/n ? [n] :**y**

Image synchronization is in progress, please wait.

Installing Loader, please wait.

Installing Loader on module 5 ... successful

Installing Loader on module 6 ... successful

Installing BIOS, please wait.

Installing BIOS on module 1 ... not required (same version)

Installing BIOS on module 2 ... not required (same version)

Installing BIOS on module 3 ... not required (same version)

Installing BIOS on module 4 ... not required (same version)

Installing BIOS on module 5 ... not required (same version)

Installing BIOS on module 6 ... not required (same version)

Installing BIOS on module 7 ... not required (same version)

Installing BIOS on module 8 ... not required (same version)

Installing BIOS on module 9 ... not required (same version)

Updating boot variables .. successful

Saving configuration, please wait.

Reload of the standby supervisor is in progress, please wait

Success, the standby supervisor is online and ready to takeover

Send documentation comments to mdsfeedback-doc@cisco.com

Example 5-3 displays the file output on the console of the standby supervisor module. It is a continuation of the output displayed in Example 5-1 and Example 5-2.

Example 5-3 Output of the install all Command on the Standby Supervisor Module Console

```
Installation procedure in progress, please wait.
The login will be disabled until the installation is completed.
Switchover to this supervisor is successful
Install of module 1 is in progress, please wait.
Install of module 2 is in progress, please wait.
Install of module 3 is in progress, please wait.
Install of module 4 is in progress, please wait.
Install of module 7 is in progress, please wait.
Install of module 8 is in progress, please wait.
Install of module 9 is in progress, please wait.
```

The installation procedure has completed successfully.

Example 5-4 displays the **install all** command output of a TFTP GET operation.

Example 5-4 Successful TFTP Get Operation

```
switch# install all system tftp://10.24.3.36/m9500-1.0.3 kickstart tftp
://10.24.3.36/m9500-kickstart.1.0.3
Copying tftp://10.24.3.36/m9500-1.0.3 to bootflash:/m9500-1.0.3
Trying to connect to tftp server.....
*****
TFTP get operation was successful

Copying tftp://10.24.3.36/m9500-kickstart.1.0.3 to bootflash:/m9500-kickstart.1.0.3
Trying to connect to tftp server.....
*****
TFTP get operation was successful

Image verification is in progress, please wait.
Jan 29 05:10:51 switch %LOCAL2-1-SYSTEM_MSG:      root : command not allowed
; TTY=unknown ; PWD=/var/tmp/4683 ; USER=root ; COMMAND=mount -o ro /mnt/bootloa
der
Jan 29 05:10:51 switch %LOCAL2-1-SYSTEM_MSG:      root : command not allowed
; TTY=unknown ; PWD=/var/tmp/4683 ; USER=root ; COMMAND=umount /mnt/bootloader
This command is going to install system image m9500-sflek9-mz.1.0.3.bin
and kickstart image m9500-sflek9-kickstart-mz.1.0.3.upg.bin
on this system
The command will:
- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the standby supervisor
- Perform a HA Switchover
- Perform a hitless upgrade of module 1, 9

Do you want to continue y/n ? [n] : y

Image synchronization is in progress, please wait.

Installing Loader, please wait.
Installing Loader on module 5 ... successful
Installing Loader on module 6 ... successful

Installing BIOS, please wait.
Installing BIOS on module 1 ... not required (same version)
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
Installing BIOS on module 5 ... not required (same version)
Installing BIOS on module 6 ... not required (same version)
Installing BIOS on module 9 ... not required (same version)

Updating boot variables .. successful
Saving configuration, please wait.
Reload of the standby supervisor is in progress, please wait
Success, the standby supervisor is online and ready to takeover
Synchronizing with Firmbase...
```

Example 5-5 displays the **install all** command output of a failed operation due to a full standby disk.

Example 5-5 Failed Operation due to a Full Standby bootflash: Directory

```
switch# install all system tftp://10.24.3.36/m9500-1.0.3 kickstart tftp:
//10.24.3.36/m9500-kickstart.1.0.3
Copying tftp://10.24.3.36/m9500-1.0.3 to bootflash:/m9500-1.0.3
Trying to connect to tftp server.....
*****
TFTP get operation was successful

Copying tftp://10.24.3.36/m9500-kickstart.1.0.3 to bootflash:/m9500-kickstart.1.0.3
Trying to connect to tftp server.....
*****
TFTP get operation was successful

Image verification is in progress, please wait.
Jan 28 12:14:24 switch %LOCAL2-1-SYSTEM_MSG:      root : command not allowed
; TTY=unknown ; PWD=/var/tmp/2990 ; USER=root ; COMMAND=mount -o ro /mnt/bootloader
Jan 28 12:14:24 switch %LOCAL2-1-SYSTEM_MSG:      root : command not allowed
; TTY=unknown ; PWD=/var/tmp/2990 ; USER=root ; COMMAND=umount /mnt/bootloader
This command is going to install system image m9500-sflek9-mz.1.0.3.bin
and kickstart image m9500-sflek9-kickstart-mz.1.0.3.upg.bin
on this system
The command will:
- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the standby supervisor
- Perform a HA Switchover
- Perform a hitless upgrade of module 1, 9

Do you want to continue y/n ? [n] : y

Image synchronization is in progress, please wait.
Image Synchronization failed(Standby disk may be full) <-----insufficient space message
```

Example 5-6 displays the **install all** command output of a failed operation due to an invalid image.

Example 5-6 Failed Operation due to an Invalid Image

```
switch# install all system kickstart-1.0.3 kickstart system-1.0.3
The system image is not a valid image
System Image type-check did not succeed
```

Send documentation comments to mdsfeedback-doc@cisco.com

Performing a Manual Upgrade on a Dual Supervisor Switch



Caution

If you are a new user, use the **install all** command to perform a software upgrade. The information and instructions provided in this section targets administrators or individuals who are completely familiar with specific switch functions.

You can manually upgrade any switch in the Cisco MDS 9200 Series or the Cisco MDS 9500 Series using the procedures provided in this section. This upgrade mechanism requires you to implement some or all procedures depending on your switch or network configuration.

To perform a manual upgrade for a dual supervisor switch, follow these steps.

-
- Step 1** [“Preparing for a Manual Installation” section on page 5-12](#)
 - Step 2** [“Upgrading a Loader” section on page 5-15](#)
 - Step 3** [“Upgrading the BIOS” section on page 5-16](#)
 - Step 4** [“Specifying Kickstart and System Boot Variables” section on page 5-18](#)
 - Step 5** [“Performing a System Switchover” section on page 5-19](#)
-

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 Configuration” section on page 3-27](#)).
 - Step 3** Copy the software image from a TFTP location to one of two targets: bootflash: or slot0:.

The switch remains operational while the image file is copied.

- Bootflash device (TFTP defaults to the bootflash device)—Copy the software image file from the appropriate TFTP directory to bootflash.

```
switch# copy tftp://<server IP address>/<file name in TFTP> <destination file name as desired>
```

For example:

```
switch# copy tftp://10.1.7.2/system_image bootflash:system_image
```



Note

The Cisco MDS 9200 Series of switches do not have an external CompactFlash (see the [“About Flash Devices” section on page 3-22](#)). If you are using a switch in this series, use the bootflash: directory to copy and verify files.

- CompactFlash device—Copy the software image file from the appropriate TFTP directory to the CompactFlash device in slot0:.

```
switch# copy tftp://<server IP address>/<file name in TFTP> slot0:system_image
```

Send documentation comments to mdsfeedback-doc@cisco.com

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: partition, 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:
admin      524288   Jan 22 04:27:42 1980 additionalfile1
admin      14576128  Jan 07 20:56:58 1980 additionalfile2
admin      3829262   Jan 07 20:50:58 1980 additionalfile3
admin      18961190   Jan 07 22:30:34 1980 system_image
admin      14514688   Jan 07 22:27:21 1980 kickstart_image
root       12288     Dec 22 23:32:06 2002 lost+found/
admin      2373      Jan 04 23:35:12 1980 z1.cfg
          Usage for bootflash: filesystem
          114538496 bytes total used
          61939712 bytes free
          186085376 bytes available
```

Step 5 Ensure that the software images are not damaged or corrupted in the saved bootflash: location.

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-sflek9-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 version compatibility** command.

```
switch# show version compatibility bootflash:system_image
```

Step 7 Verify if the standby supervisor module has sufficient space for the new image files.

a. Attach to the standby module to verify the space requirements.

```
switch# attach module 6
Attaching to module 6 ...
To exit type 'exit', to abort type '$.'
```

b. Change to the bootflash: directory.

```
switch(standby)# cd bootflash:
```

c. Verify if the standby supervisor module has sufficient space for the new image files.

```
switch(standby)# dir

admin      14457344   Jan 29 06:10:18 1980 kickstart-1.0.3
admin      14584832   Jan 30 02:10:02 1980 kickstart_image3
admin      14584832   Jan 31 08:25:21 1980 kickstart_image3a
admin      12288      Jan 29 00:25:33 1980 lost+found/
admin      19189547   Jan 29 06:10:01 1980 system-1.0.3
admin      19200399   Jan 30 02:09:45 1980 system_image3
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

admin    19153046   Jan 31 08:25:05 1980 system_image3a
Usage for bootflash: filesystem
105803776 bytes total used
70674432 bytes free
186085376 bytes available

```

- d. Delete files, if required, to make more space for the new image files.

```

switch(standby)# del kickstart-1.0.3
switch(standby)# del system-1.0.3

```

- e. Exit the standby supervisor module prompt.

```

switch(standby)# exit
rlogin: connection closed.

```

Step 8 Ensure that the auto-sync feature is disabled.

- a. Use the **show auto-sync** command in EXEC mode to verify if this option is configured.

```

switch# show system auto-sync
auto-sync is disabled
auto-sync not started

```

If the **system auto-sync** command is disabled (default), skip to [Step 9](#).

If **system auto-sync** command is enabled, first disable this option by continuing with [Step 8b](#).

- b. Change to configuration mode.

```

switch# config terminal

```

- c. Disable the **auto-sync** option on the active supervisor module.

```

switch(config)# no system auto-sync image

```

- d. Exit to the EXEC mode.

```

switch(config)# exit

```

Step 9 Copy the system and kickstart image files to the standby supervisor module. The following example assumes the standby supervisor module is in slot 6:

```

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

```

Step 10 Use the **show version image** command to verify the integrity of the image before loading the images. This command can be used for both the system and kickstart images.

```

switch(boot)# show version image bootflash:system_image <-----system image
image name: m9500-sflek9-mz.1.0.3.bin
bios:      version   v1.0.6(01/27/03)
system:    version 1.0(3)
compiled:  01/25/2003 12:00:00

switch(boot)# show version image bootflash:kickstart_image <-----kickstart image
image name: m9500-sflek9-kickstart-mz.1.0.3.upg.bin
kickstart: version 1.0(3)
loader:    version 1.0(3)
compiled:  01/25/2003 12:00:00

switch# show version image bootflash:bad_image <-----failure case
Md5 Verification Failed
Image integrity check failed

```


Send documentation comments to mdsfeedback-doc@cisco.com

Upgrading a Loader

Using the **install module slot# of the supervisor module loader** command, you can upgrade 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 will not be impacted.



Caution

Before issuing this command, be sure to read the release notes to verify compatibility issues between the boot 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 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.
Software
  BIOS:      version 1.0(3)
  loader:    version 1.0(2) <----- current running version
  kickstart: version 1.0(3)
  system:    version 1.0(3)

  BIOS compile time:      01/27/03
  kickstart image file is: bootflash:/kickstart_imageu
  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** 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 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.
Software
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

BIOS:          version 1.0(3)
loader:        version 1.0(3) [last : 1.0(3)] <-----new running version
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
```

```

bootflash: 503808 blocks (block size 512b)
slot0:     500736 blocks (block size 512b)

```

```
switch uptime is 0 days 3 hours 56 minute(s) 6 second(s)
```

Last reset

```

Reason: Watchdog Timeout/External Reset
System version: 1.0(3)

```

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 or switching module BIOS only if a new BIOS image is provided by Cisco. 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 effective. You can schedule the reboot at a convenient time so traffic will not be 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
Andiamo 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)

```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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 detail
image name: m9500-sflek9-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 version 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
###
BIOS upgrade succeeded for module 1

```



Caution

Do not reboot the switch if any errors were indicated in response to this command.

Step 4 Connect to the switching module using the **attach module** command.

```

switch# attach module 6
switch(standby)#

```

Step 5 Issue the **show version** command to verify that the module was updated with a the new BIOS version.

```

switch(standby)# show version
Software
bios:      version  v1.0.6(01/27/03) <----- New BIOS same version
system:    version  1.0(3)
compiled:   2/28/2003 5:00:00
...

```

Send documentation comments to mdsfeedback-doc@cisco.com

Specifying Kickstart and System Boot Variables

To specify the kickstart and system images in one or both supervisor modules (active and standby), follow these steps:

Step 1 Log into the switch through the console port, an SSH session, or a Telnet session.

Step 2 Change to configuration mode.

```
switch# config terminal
```

Step 3 Specify the kickstart image to be used for the reboot.

```
switch(config)# boot kickstart bootflash:kickstart_image
```



Note You can only specify one image for the KICKSTART variable.

Step 4 Specify the first image (system image).

```
switch(config)# boot system bootflash:system_image
```

Step 5 Verify that the Flash device is physically in slot0: before issuing the next command.

Step 6 Specify the second image (kickstart image) for the reboot.

```
switch(config)# boot system slot0:kickstart_image
```

Step 7 Change to the EXEC mode.

```
switch(config)# exit
```

Step 8 Issue the **show boot** command, to display the current contents of the SYSTEM variable.

```
switch# show boot
sup-1
KICKSTART variable = bootflash:/kickstart_image
SYSTEM variable = bootflash:/system_image;
sup-2
KICKSTART variable = bootflash:/kickstart_image
SYSTEM variable = bootflash:/system_image;
```

To clear the current contents of the SYSTEM variable in both supervisor modules, use the **no boot** command.

```
switch(config)# no boot system
switch# show boot
sup-1
KICKSTART variable = bootflash:/kickstart_image
SYSTEM variable not set
sup-2
KICKSTART variable = bootflash:/kickstart_image
SYSTEM variable not set
```

To clear the current contents of the SYSTEM variable in one supervisor module, use the **sup-1** or **sup-2** options:

```
switch(config)# no boot system sup-2
switch# show boot
sup-1
KICKSTART variable = bootflash:/kickstart_image
SYSTEM variable = bootflash:/system_image
sup-2
KICKSTART variable = bootflash:/kickstart_image
```

Send documentation comments to mdsfeedback-doc@cisco.com

SYSTEM variable not set

- Step 9** Save the new variable configuration so the new image is used the next time you log into the switch.

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

- Step 10** Reload the standby supervisor module to verify the effect of the new image on the supervisor module and the switching modules in the switch.

```
switch# reload module 6
This command will reboot the system. (y/n)? y
```

The **reload** command reboots the system and updates the variable in one or both supervisor modules automatically. Use the **reload** command after the configuration information is entered into a file and saved to the startup configuration.

- Step 11** Issue the **show module** command to verify that the supervisor module in Slot 6 is running the new image.

```
switch# sh module
```

| Mod | Ports | Module-Type | Model | Status |
|-----|-------|---------------------|-----------------|------------|
| 1 | 16 | 1/2 Gbps FC Module | DS-X9016 | ok |
| 5 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | active * |
| 6 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | ha-standby |
| 9 | 16 | 1/2 Gbps FC Module | DS-X9016 | ok |

| Mod | Sw | Hw | World-Wide-Name(s) (WWN) |
|-----|--------|-------|--|
| 1 | 1.0(3) | 0.3 | 20:01:00:05:30:00:13:9e to 20:10:00:05:30:00:13:9e |
| 5 | 1.0(3) | 0.602 | -- |
| 6 | 1.0(3) | 0.0 | -- <----- New running version in module 6 |
| 9 | 1.0(3) | 0.0 | 22:01:00:05:30:00:13:9e to 22:10:00:05:30:00:13:9e |

| Mod | MAC-Address(es) | Serial-Num |
|-----|--|-------------|
| 1 | 00-05-30-00-81-6e to 00-05-30-00-81-72 | jab063908gj |
| 5 | 00-05-30-00-84-1a to 00-05-30-00-84-1e | jab063909cv |
| 6 | 00-05-30-00-2c-5e to 00-05-30-00-2c-62 | |
| 9 | 00-05-30-00-03-0c to 00-05-30-00-03-10 | 123 |

· this terminal session



Note

The next time you reboot the switch, the saved image is used. If you do not save the configuration, the previously saved startup configuration image is used.

Performing a System Switchover

A switch in the Cisco MDS 9500 Series has two supervisor modules. For example a Cisco MDS 9509 director has one in slot 5 (sup-1) and one in slot 6 (sup-2). When both supervisor modules power up at the same time, the module in slot 5 enters the active mode, while the second module in slot 6 enters the standby mode.

This procedure assumes that slot 5 contains the currently active supervisor module and slot 6 contains the standby supervisor module.

Send documentation comments to mdsfeedback-doc@cisco.com

**Tip**

Verify that the standby supervisor module has sufficient space to perform this procedure.

To perform a system switchover on a dual supervisor module, follow these steps:

- Step 1** Switch over to the standby supervisor module (in slot 6) which becomes the active module.

```
switch# system switchover
```

The newly-active supervisor module in slot 6 takes over as the active supervisor.

**Note**

The previously-active supervisor module in slot 5 reboots automatically.

All incompatible switching modules are automatically detected and upgraded to the same image version as the supervisor modules. All compatible modules must be manually upgraded to be the same version as the supervisor modules.

If you are on a console port, you will see the rebooting messages as the supervisor modules come up.

**Note**

If you are on a Telnet or SSH session, the session is terminated.

- Step 2** Use the **show module** command to determine if a switching module must be updated.

If the switching module is running an image version that is older than the image version on the supervisor module, that switching module must be updated using this procedure.

- Step 3** Issue the **install module slot# image** command on the affected module.

```
switch# install module 8 image
SW version of linecard image in Supervisor = 1.0(3)
Module 8 : SW version = 1.0(3) , Image version is same, upgrade is not required
```

The preceding example displays a module that does not need to be updated. The following example displays a module that requires an update.

```
switch# install mod 9 image
SW version of linecard image in Supervisor = 1.1(1)
Module 9 : SW version = 1.0(3) , Non disruptive image upgrade is possible
Starting non disruptive upgrade of module 9
Feb 3 02:37:13 switch %LC-2-MSG:SLOT9 LC-2-IMG_DNLD_STARTED: Module image
download process. Please wait until completion...
Feb 3 02:37:33 switch %LC-2-MSG:SLOT9 LC-2-IMG_DNLD_COMPLETE: Module image
download process. Download successful.
Module 9 upgrade successful
```

This command updates the specified switching module using the image version that is running on the active supervisor module.

Repeat this step for each switching module that must be updated.

- Step 4** Issue the **show module** to verify the module version and the **show system redundancy status** command to ensure the modules are functioning as required.

Send documentation comments to mdsfeedback-doc@cisco.com

Performing a Manual Upgrade on a Single Supervisor Switch

This procedure can be performed on any switch in the Cisco MDS 9200 Series. While it can also be performed on any Cisco MDS 9500 Series switch that has only one supervisor module, this is not recommended.

To manually upgrade a switch with a single supervisor module, follow these steps:

Step 1 Update the boot variables on the supervisor modules (see the “[Specifying Kickstart and System Boot Variables](#)” section on page 5-18).

Step 2 Issue the **show version compatibility** command to view the effects of an installation.

```
switch# show version compatibility bootflash:system_image
Version comparison between /bootflash/system_image and running-image:
Mod No   Mod Type      SRG Compare Result
1         FC/SUP        Non-Disruptive upgrade is possible
2         LC           Linecard version is compatible
```

Step 3 Copy the software image from a remote location to the bootflash: directory.

```
switch# copy scp://171.00.16.26//tftpboot/NDrel/qa/excal/final/kickstart_image
bootflash:kickstart_image
Enter username:user
user@171.00.16.26's password:
m9200-ek9-kickstart- 100% |*****| 14219 KB      00:16
```

Step 4 Verify that the file was copied in the required directory.

```
switch# dir bootflash:

root          1024   Dec 31 16:04:53 1979 .ssh/
root          1024   Dec 31 16:04:53 1979 .ssh2/
admin         524288  Jan 21 15:31:51 2003 MC1118.BIN
admin        14552064  Feb 05 19:27:11 1980 kick-1.0.2.33a
admin        14560256  Feb 14 11:32:30 1980 kick-1.0.2.37a
admin        14560256  Feb 14 11:16:56 1980 kick-1.0.2.37au
root          12288   Dec 31 16:04:42 1979 lost+found/
admin        19003669  Feb 05 19:28:11 1980 system-1.0.2.33a
admin        18967721  Feb 14 11:31:18 1980 system-1.0.2.37a
admin        18927679  Feb 14 11:19:20 1980 system-1.0.2.37au

Usage for bootflash: filesystem
105733120 bytes total used
70745088 bytes free
186085376 bytes available
```

Step 5 Ensure that the software images are not damaged or corrupted in the saved bootflash: location.

When copying a new image to your switch, you should 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: m9200-ek9-kickstart
kickstart:  version 1.0(3)
loader:      version 1.0(3)
compiled:    2/12/2003 11:00:00
```

Send documentation comments to mdsfeedback-doc@cisco.com

**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 Issue the **install all** command on the supervisor module.

**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.

```
switch# install all system bootflash:system_image kickstart bootflash:kickstart_image
```

Image verification is in progress, please wait.

This command is going to install system image m9200-ek9-mz.1.0.2.37a.bin
and kickstart image m9200-ek9-kickstart-mz.1.0.2.37a.bin
on this system

The command will:

- Install the Loader, if required
- Install the BIOS, if required
- Update boot variables
- Save configuration
- Reload the supervisor

Do you want to continue y/n ? [n] : **y**

Installing Loader, please wait.

Installing Loader on module 1 ... not required (same version)

Installing BIOS, please wait.

Installing BIOS on module 1 ... not required (same version)

Updating boot variables .. successful

Saving configuration, please wait.

Synchronizing with Firmware...

General Software Firmware[r] SMM Kernel 1.1.1002 Oct 11 2002 13:36:57

Copyright (C) 2002 General Software, Inc.

Firmware initialized.

...

00000589K Low Memory Passed

00979840K Ext Memory Passed

Wait.....

...

General Software Pentium III Embedded BIOS 2000 (tm) Revision 1.0.(6)

(C) 2002 General Software, Inc.ware, Inc.

Pentium III-1.0-6E69-AA6E

```
+-----+
|           System BIOS Configuration, (C) 2002 General Software, Inc.           |
+-----+
| System CPU           : Pentium III       | Low Memory           : 630KB       |
| Coprocessor          : Enabled           | Extended Memory        : 957MB       |
| Embedded BIOS Date   : 01/27/03         | ROM Shadowing          : Enabled     |
+-----+
```

Loader Loading stage1.5.

Loader loading, please wait...

...

Auto booting bootflash:/kick-1.0.2.37a bootflash:/system-1.0.2.37a;...

Send documentation comments to mdsfeedback-doc@cisco.com

```
Booting kickstart image: bootflash:/kick-1.0.2.37a...
.....Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
Uncompressing system image: bootflash:/system-1.0.2.37a
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
Uncompressing linecard components
INIT: Entering runlevel: 3

MDS Switch
switch login: admin
Password:
```

Step 7 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 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.

Software
  BIOS:      version 1.0.6
  loader:    version 1.0(2.37a)
  kickstart: version 1.0(3) [build 1.0(2.37a)]
  system:    version 1.0(3) [build 1.0(2.37a)]

  BIOS compile time:      01/27/03
  kickstart image file is: bootflash:/kick-1.0.2.37a
  kickstart compile time: 3/6/2003 12:00:00
  system image file is:   bootflash:/system-1.0.2.37a
  system compile time:    3/6/2003 12:00:00

Hardware
  RAM 963052 kB
  bootflash: 503808 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)
  switch uptime is 0 days 0 hour 2 minute(s) 36 second(s)
  Last reset at 433817 usecs after Thu Feb 14 11:34:31 1980
    Reason: Reset Requested by CLI command reload
  System version: 1.0(4v)
```

Step 8 Use the **show module** command to determine if the switching module must be updated.

If the switching module is running an image version that is older than the image version on the supervisor module, that switching module must be updated. If the switching module image is not compatible with the supervisor module image, the switching module image will be automatically updated.

Step 9 Issue the **install module slot# image** command on the affected module. The following example displays a module that does not need to be updated.

```
switch# install module 2 image
SW version of linecard image in Supervisor = 1.0(4u)
Module 2 : SW version = 1.0(4u) , Image version is same, upgrade is not required
```

The following example displays a module that needs to be updated.

```
switch# install mod 2 image
SW version of linecard image in Supervisor = 1.0(3)
Module 2 : SW version = 1.0(4u) , Non disruptive image upgrade is possible
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
Starting non disruptive upgrade of module 2
Feb  3 02:37:13 switch %LC-2-MSG:SLOT2 LC-2-IMG_DNLD_STARTED:  Module image
download process. Please wait until completion...
Feb  3 02:37:33 switch %LC-2-MSG:SLOT2 LC-2-IMG_DNLD_COMPLETE:  Module image
download process. Download successful.
Module 2 upgrade successful
```

This command updates the specified switching module using the image version that is running on the active supervisor module.

Quick, One-Step Upgrade

To perform a quick, one-step 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 Files”](#) section on page 3-28).
 - Step 2** Set the boot variables (see the [“Specifying Kickstart and System Boot Variables”](#) section on page 5-18).
 - Step 3** Issue the **reload** command. The **reload** command reboots the system. This upgrade is disruptive.
-

Send documentation comments to mdsfeedback-doc@cisco.com

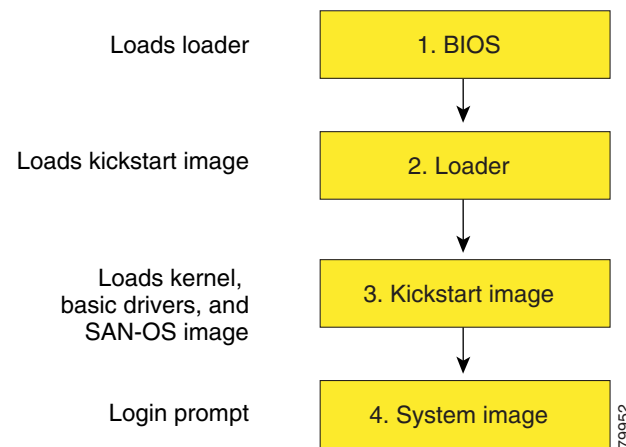
Recovering a Corrupted Bootflash

All switch configurations reside in the internal bootflash. If you have a corrupted internal bootflash, you could potentially lose your configuration. Be sure to save and back up your configuration files periodically.

The regular switch boot goes through the following sequence (see [Figure 5-2](#)):

1. The basic input/output system (BIOS) loads the boot loader.
2. The boot loader loads the kickstart image into RAM and starts the kickstart image.
3. The kickstart image loads and starts the system image.
4. The system image reads the startup configuration file.

Figure 5-2 Regular Boot Sequence



If the images on your switch are corrupted and you are not able to proceed (error state), you can determine the reason and attempt to interrupt the switch boot sequence and recover the image by entering the BIOS configuration utility described in the following section. Access this utility only when needed to recover a corrupted internal disk.



Caution

The BIOS changes explained in this section are required only if you need to recover a corrupted bootflash.

Send documentation comments to mdsfeedback-doc@cisco.com

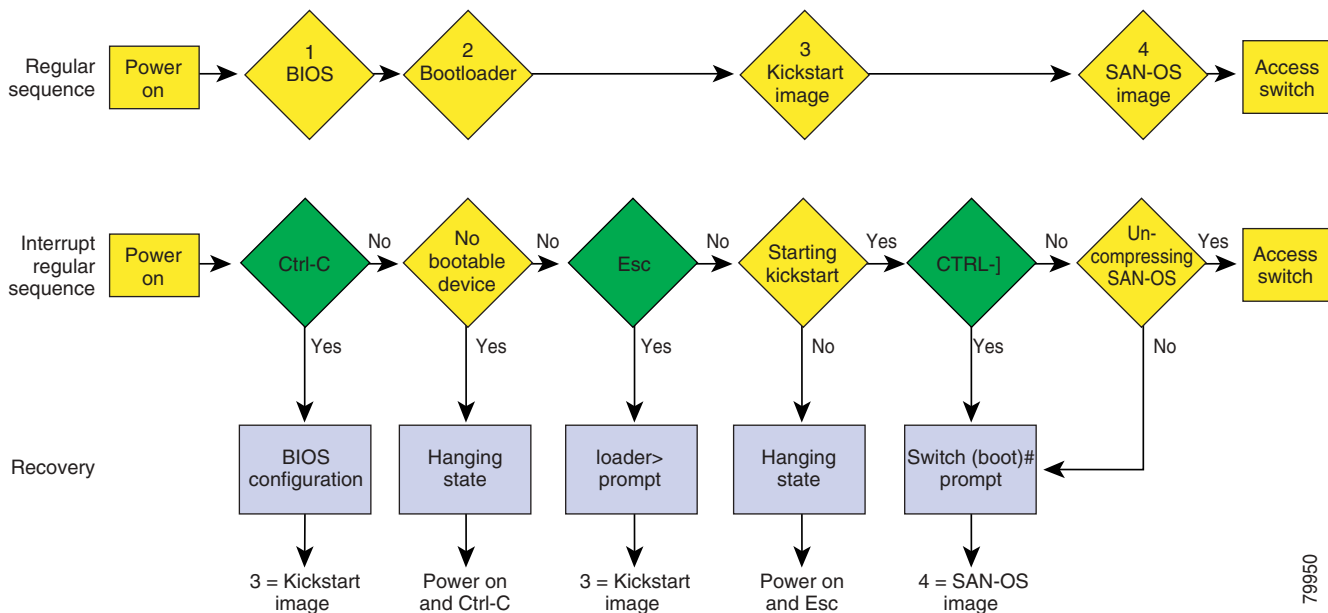
Recovery procedures require the regular sequence to be interrupted. The internal switch sequence goes through four phases between the time you turn the switch on and the time the switch prompt appears on your terminal—BIOS, boot loader, Kickstart, and system (see [Table 5-1](#) and [Figure 5-3](#)).

Table 5-1 Recovery Interruption

| Phase | Normal Prompt ¹ | Recovery Prompt ² | Description |
|-------------|----------------------------|------------------------------|--|
| BIOS | loader> | No bootable device | The BIOS begins the power-on self test, memory test, and other operating system applications. While the test is in progress, press Ctrl-C to enter the BIOS configuration utility and use the netboot option. |
| Boot loader | Starting kickstart | loader> | The boot loader uncompresses loaded software to boot an image using its file name as reference. These images are made available through bootflash. When the memory test is over, press Esc to enter the boot loader prompt. |
| Kickstart | Uncompressing system | switch (boot) # | When the boot loader phase is over, press Ctrl-] (Control key plus right bracket key) to enter the <code>switch (boot) #</code> prompt. If the corruption causes the console to stop at this prompt, copy the system image and reboot the switch. |
| System | Login: | — | The system image loads the configuration file of the last saved running configuration and returns a switch login prompt. |

1. This prompt or message appears at the end of each phase.
2. This prompt or message appears when the switch cannot progress to the next phase.

Figure 5-3 Regular and Recovery Sequence



79950

Send documentation comments to mdsfeedback-doc@cisco.com

Recovery Using BIOS Setup

To recover a corrupted bootflash image (no bootable device found message) for a switch with a single supervisor module, follow these steps:

- Step 1** Connect to the console port of the required switch.
- Step 2** Boot or reboot the switch.
- Step 3** Press **Ctrl-C** to interrupt the BIOS setup during the BIOS memory test.
You see the netboot BIOS Setup Utility screen (see [Figure 5-4](#)).

Figure 5-4 BIOS Setup Utility



Note

Your navigating options are provided at the bottom of the screen.

Tab = Jump to next field

Ctrl-E = Down arrow

Ctrl-X = Up arrow

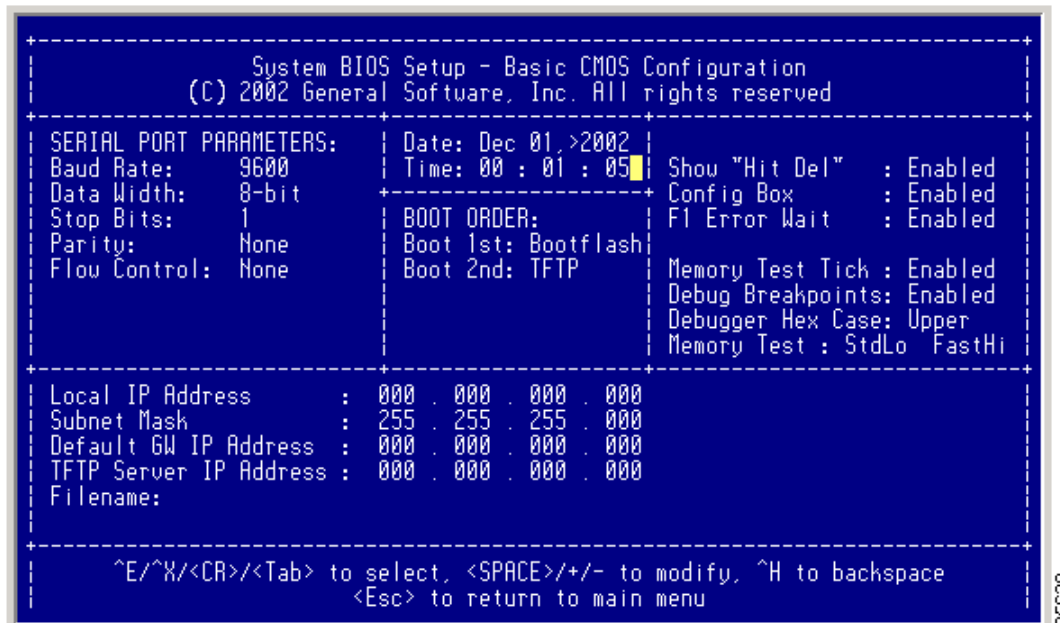
Ctrl-H = Erase (Backspace might not work if your terminal is not configured properly.)

Send documentation comments to mdsfeedback-doc@cisco.com

Step 4 Press the **Tab** key to select the Basic CMOS Configuration, and press **Enter**.

You see the BIOS setup CMOS Configuration screen (see Figure 5-5).

Figure 5-5 BIOS Setup Configuration (CMOS)



Step 5 Change the "Boot 1st:" field to **TFTP**.

Step 6 Press the **Tab** key until you reach the local IP Address field.

Step 7 Enter the local IP address for the switch, and press the **Tab** key.

Step 8 Enter the subnet mask for the IP address, and press the **Tab** key.

Step 9 Enter the IP address of the default gateway, and press the **Tab** key.

Step 10 Enter the IP address of the TFTP server, and press the **Tab** key.

Step 11 Enter the image name (kickstart), and press the **Tab** key. This path should be relative to the TFTP server root directory.



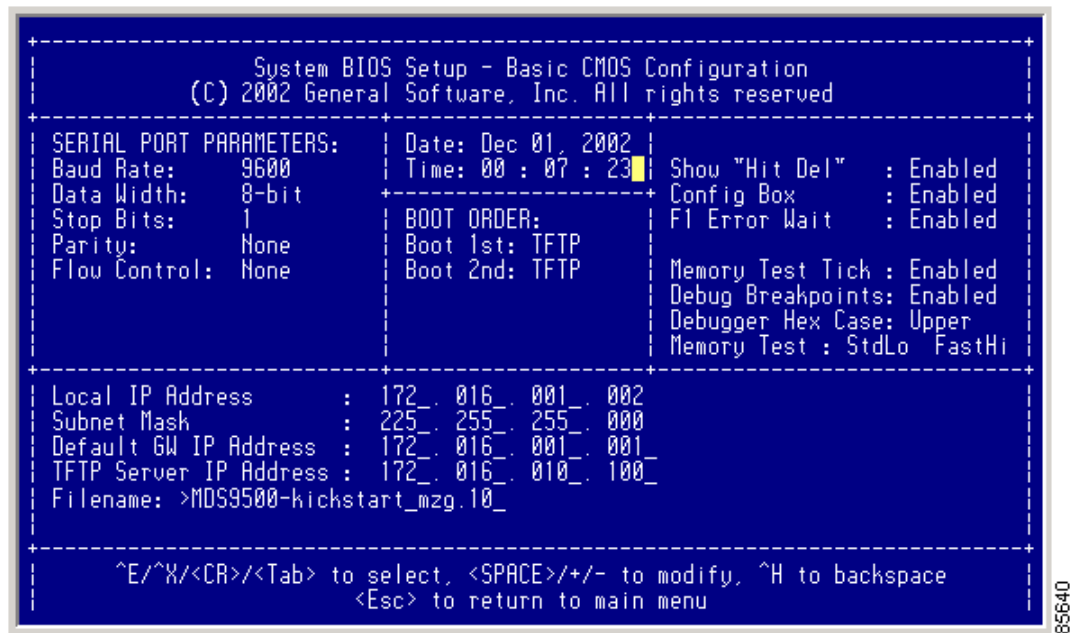
Caution

The file name must be entered exactly as it is displayed on your TFTP server. For example, if you have a file name **MDS9500-kiskstart_mzg.10**, then enter this name using the exact upper case characters and file extensions as shown on your TFTP server.

Send documentation comments to mdsfeedback-doc@cisco.com

You see the configured changes (see [Figure 5-6](#)).

Figure 5-6 BIOS Setup Configuration (CMOS) Changes



Step 12 Press the **Esc** key to return to the main menu.

Step 13 Choose **Write to CMOS and Exit** from the main screen to save your changes.



Note These changes are saved in the CMOS.



Caution The switch must have IP connectivity to reboot using the newly configured values.

You are placed at the following prompt:

```
switch(boot) #
```

Step 14 Enter the **init system** command at the `switch(boot)` prompt, and press **Enter**.

```
switch(boot) # init system
```

The `switch(boot) #` prompt indicates that you have a usable kickstart image.



Note The **init system** command also installs a new loader from the existing (running) kickstart image.

Step 15 Follow the procedure specified in the “[Recovery from the switch\(boot\)# Prompt](#)” section on page 5-31.

Send documentation comments to mdsfeedback-doc@cisco.com

Recovery from the loader> Prompt

To recover a corrupted kickstart image (system error state) for a switch with a single supervisor module, follow these steps:

- Step 1** Press the **Esc** key to interrupt the boot loader setup after the BIOS memory test.



Note Press **Esc** immediately after you see the following message:

```
00000589K Low Memory Passed
00000000K Ext Memory Passed
Hit ^C if you want to run SETUP....
Wait.....
```

If you wait too long, you will skip the boot loader phase and enter the kickstart phase.

You see the loader> prompt.



Caution The loader> prompt is different from the regular switch# or switch(boot)# prompt. The CLI command completion feature does not work at this prompt and may result in undesired errors. You must type the command exactly as you want the command to appear.

- Step 2** Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
loader> ip address 172.16.1.2 255.255.255.0
Found Intel EtherExpressPro100 82559ER at 0xe800, ROM address 0xc000
Probing...[Intel EtherExpressPro100 82559ER]Ethernet addr: 00:05:30:00:52:27
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 0.0.0.0
```

- Step 3** Enter the IP address of the default gateway, and press **Enter**.

```
loader> ip default-gateway 172.16.1.1
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 172.16.1.1
```

- Step 4** Boot the kickstart image file from the required TFTP server, and press **Enter**.

```
loader> boot tftp://172.16.10.100/kickstart-latest
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Bootimg: /kick-282 console=ttyS0,9600n8nn quiet loader_ver="1.0(2)"....
.....Image verification OK
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
WARNING: image sync is going to be disabled after a loader netboot
Loading system software
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
INIT: Sending processes the TERM signal
switch(boot)#
```

The `switch(boot)#` prompt indicates that you have a usable Kickstart image.

Step 5 Copy the system and kickstart images again.

```
switch(boot)# copy tftp://172.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server.....
```

```
switch(boot)# copy tftp://172.16.10.100/system-img bootflash:kickstart-img
Trying to connect to tftp server.....
```

Step 6 Follow the procedure specified in the “[Recovery from the switch\(boot\)# Prompt](#)” section on page 5-31.

Recovery from the switch(boot)# Prompt

To recover a system image using the kickstart image for a switch with a single supervisor module, follow these steps:

Step 1 Follow this step if you issued a `init system` command. Otherwise, skip to [Step 2](#).

a. Change to configuration mode and configure the IP address of the switch’s `mgmt0` interface.

```
switch(boot)# config t
switch(boot) (config)# interface mgmt0
```

b. Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip address 172.16.1.2 255.255.255.0
```

Step 2 Issue the `no shut` command to enable the interface on the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# no shut
```

Step 3 Follow this step if you issued a `init system` command. Otherwise, skip to [Step 4](#).

a. Enter the IP address of the default gateway, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip default-gateway 172.16.1.1
```

Step 4 Exit to EXEC mode.

```
switch(boot) (config-mgmt0)# end
```

Step 5 Copy the system image from the required TFTP server, and press **Enter**.

```
switch(boot)# copy tftp://172.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server.....
```

Step 6 Copy the kickstart image from the required TFTP server, and press the **Enter** key.

```
switch(boot)# copy tftp://172.16.10.100/system-img bootflash:kickstart-img
Trying to connect to tftp server.....
```

Step 7 Verify that the system and kickstart image files are copied to your bootflash: directory.

```
switch(boot)# dir bootflash:
root          524288  Dec 31 16:07:52 1979  additionalfile1
admin        13777408  Feb 13 07:20:07 2003  kickstart_image
admin         1199287  Feb 12 10:12:43 2003  debug-plugin
admin           104   Jan 20 11:49:56 2003  flt
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
admin      24991890  Feb 13 07:20:37 2003 system_image
root       12288    Dec 10 18:23:47 2002 lost+found/
root       43      Dec 24 15:02:57 2002 test.flt
```

```
Usage for bootflash: filesystem
      84705280 bytes total used
      325267456 bytes free
      432293888 bytes available
```

Step 8 Load the system image from the bootflash: directory:

```
switch boot# load bootflash:system_image
Uncompressing system image: bootflash:/system_image
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
MDS Switch
Would you like to enter the initial configuration mode?(yes/no): yes
```

See the “Initial Setup Routine” section on page 3-2.



Note

If you enter **no** at this point, you will return to the `switch#` login prompt, and you must manually configure the switch.

Recognizing Error States

If you see the error messages displayed in [Figure 5-7](#) or [Figure 5-8](#), follow the procedure specified in the “Recovery Using BIOS Setup” section on page 5-27.

Figure 5-7 Error State if Powered On and Ctrl-C is Entered

```
+-----+
| System BIOS Configuration, (C) 2002 General Software, Inc. |
+-----+
| System CPU       : Pentium III | Low Memory       : 630KB |
| Coprocessor     : Enabled      | Extended Memory  : 957MB |
| Embedded BIOS Date : 09/10/02 | ROM Shadowing    : Enabled |
+-----+
Boot network name is EOBC
Local IP address: 127.1.2.1

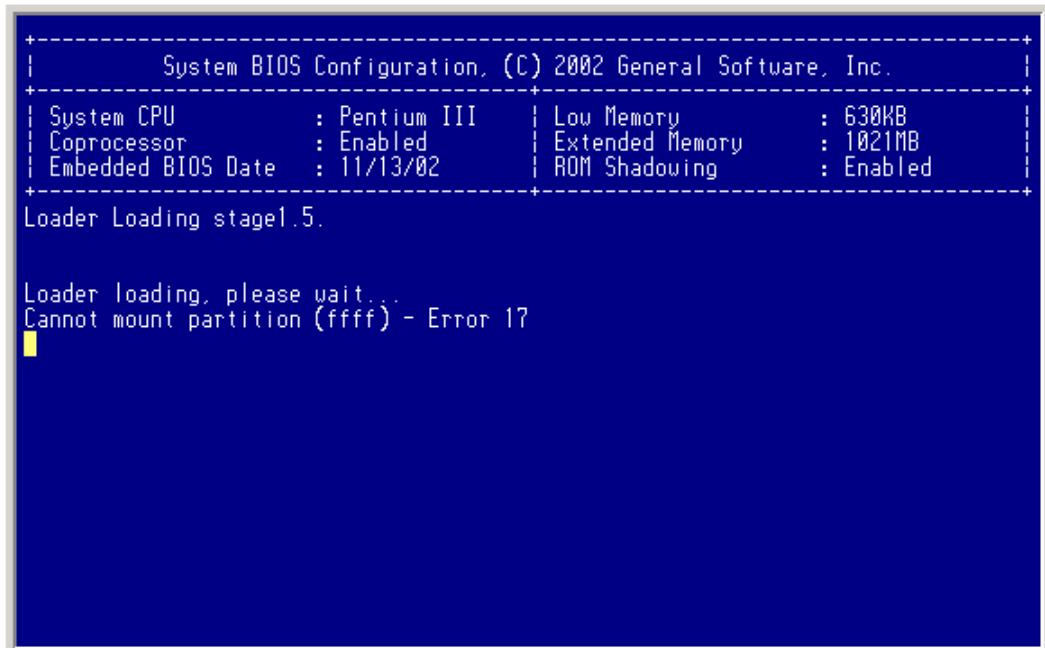
Bind to network device '/DEV/TCPIP/EOBC/BootNet'
SoBindNetName: KeOpenFile failed.
Cannot bind to the network '/DEV/TCPIP/EOBC/BootNet'
Could not get BOOTP response from the server.
BOOTNET: Dispatch duration could not be restored, reason=1.
Network boot failed, status=317.

No bootable device available.
R - REBOOT
S - SETUP
ESC - BIOS DEBUGGER
```

85642

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 5-8 Error State if Powered On and Esc is Pressed



Recovery for Switches with Dual Supervisor Modules

If one supervisor module is functioning and the other is not, boot the functioning supervisor module. Then use the booted supervisor module to bring up the supervisor module that is stuck. Issue the **reload module slot force-dnld** command (after you log into the switch) where *slot* is the slot number of the stuck supervisor module.

If both supervisor modules are not functioning, treat it like a single supervisor module recovery. First recover the image on one supervisor module and then follow the single supervisor recovery process.



Note

If the **auto-sync** option is disabled, the supervisor modules will not synchronize automatically. In this case, enable the **auto-sync** option before issuing the **reload module slot force-dnld** command. Once the synchronization is complete, disable this option.

Send documentation comments to mdsfeedback-doc@cisco.com

Replacing a Faulty Supervisor Module

The **auto-sync** option is disabled by default. If you are replacing a faulty supervisor module, ensure that the **auto-sync option** is enabled so the software versions on both modules synchronize automatically.

To replace a faulty supervisor module in a Cisco MDS 9500 Series switch, follow these steps:

-
- Step 1** Create a backup of your existing configuration file, if required (see the [“Saving the Configuration” section on page 3-27](#)).
- Step 2** Use the **show auto-sync** command in EXEC mode to verify if this option is enabled.
- If the **system auto-sync** command is disabled (default), enable this option by continuing with [Step 3](#).
- If the **system auto-sync** command is already enabled, skip to [Step 4](#).
- Step 3** Change to configuration mode and enable the **auto-sync** option.
- ```
switch# config t
switch(config)# system auto-sync image
```
- Step 4** Replace the faulty supervisor module as specified in the *Cisco MDS 9500 Series Hardware Installation Guide*.
- Step 5** Wait till the new supervisor module is online and then ensure that the synchronization was successful using the **show system auto-sync** command.
- ```
switch(config)# do show system auto-sync
administratively enabled
auto-sync successfully completed
```
- Step 6** Disable the **auto-sync** option on the active supervisor module.
- ```
switch(config)# no system auto-sync image
```
- Step 7** Exit to the EXEC mode.
- ```
switch(config)# exit
```
-

Default Factory Settings

[Table 5-2](#) lists the default settings for all Cisco MDS 9000 Family switches.

Table 5-2 Default Factory Settings

| Parameters | Default |
|------------------------|-----------------------|
| auto-sync image option | Disabled |
| Kickstart image | No image is specified |
| System image | No image is specified |



Managing Modules

This chapter describes how to manage switching modules (also known as line cards) and provides information on monitoring module states. This chapter includes the following sections:

- [About Modules, page 6-1](#)
- [Verifying the Status of a Module, page 6-2](#)
- [Viewing the State of a Module, page 6-3](#)
- [Connecting to a Module, page 6-4](#)
- [Reloading Modules, page 6-5](#)
- [Preserving Module Configuration, page 6-6](#)
- [Purging Module Configuration, page 6-7](#)
- [Powering Off Switching Modules, page 6-7](#)
- [Identifying Module LEDs, page 6-8](#)
- [Default Supervisor Module Settings, page 6-10](#)

About Modules

[Table 6-1](#) describes the supervisor module options for switches in the Cisco MDS 9000 Family.

Table 6-1 Supervisor Module Options

| Product | No. of Supervisor Modules | Supervisor Module Slot | Switching Module Features |
|----------------|--|------------------------|--|
| Cisco MDS 9216 | One module (includes 16 Fibre Channel ports) | 1 | 2-slot chassis allows one optional switching module in the other slot. |
| Cisco MDS 9509 | Two modules | 5 and 6 | 9-slot chassis allows any switching module in the other seven slots. |
| Cisco MDS 9506 | Two modules | 5 and 6 | 6-slot chassis allows any switching module in the other four slots. |

Send documentation comments to mdsfeedback-doc@cisco.com

Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.

Cisco MDS 9500 Series switches have two supervisor modules—one in slot 5 (sup-1) and one in slot 6 (sup-2). When the switch powers up and both supervisor modules come up together, the module that enters the active mode is dependent on which of the two modules 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.

Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. The switching module obtains its image from the supervisor module.

The interfaces in each module are ready to be configured when the **ok** status is displayed in the **show module** command output (see the [“Configuring Fibre Channel Interfaces”](#) section on page 9-2).

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 in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type          Model              Status
---  -
2    16     1/2 Gbps FC Module   DS-X9016           ok
5     0     Supervisor/Fabric-1   DS-X9530-SF1-K9    active *
6     0     Supervisor/Fabric-1   DS-X9530-SF1-K9    HA-standby
8    32     1/2 Gbps FC Module   DS-X9032           ok

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.0(0.253)  1.0         20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de
5    1.0(0.253)  1.0         --
6    1.0(0.253)  1.0         --
8    1.0(0.253)  1.0         20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de

Mod  MAC-Address(es)          Serial-Num
---  -
2    00-05-30-00-0f-e4 to 00-05-30-00-0f-e8  jab0636063v
5    00-05-30-00-19-66 to 00-05-30-00-19-6a  jab06370593
6    00-05-30-02-20-02 to 00-05-30-02-20-06  jab06371593
8    00-05-30-00-1a-12 to 00-05-30-00-1a-16  jab06370574
```

* 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.

Send documentation comments to mdsfeedback-doc@cisco.com



Note

A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled (see the [“HA Switchover” section on page 4-3](#)). 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 [“Viewing the State of a Module” section on page 6-3](#).

Viewing the State of a Module

If your chassis has more than one switching module (also known as line card), you will see the progress check if you issue the **show module** command several times and view the status column each time.

The switching module goes through a testing and an initializing stage before displaying an **ok** status. [Table 6-2](#) describes the possible states in which a module can exist.

Table 6-2 Module States

| show module Output | Description |
|--------------------|---|
| powered up | The hardware has electrical power. When the hardware is powered up, the software begins booting. |
| testing | The module has established connection with the supervisor and the switching module is performing bootstrap 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. In this case, issue a show environment power command to determine power consumption issues (see Chapter 26, “Monitoring System Processes and Logs”). |
| active | This module is the active supervisor module and the switch is ready to be configured. |
| HA-standby | This module is the standby supervisor module and that the HA switchover mechanism is enabled (see the “HA Switchover” section on page 4-3). |
| standby | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “HA Switchover” section on page 4-3). |

Send documentation comments to mdsfeedback-doc@cisco.com

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: <pre> 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-sflek.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 </pre> | 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 view the standby supervisor module information, but you can not configure the standby supervisor module using this command.
- On the switching module portion of the Cisco MDS 9216 supervisor module which resides in slot 1.

Send documentation comments to mdsfeedback-doc@cisco.com

Reloading Modules

You can reload the entire switch, reset specific modules in the switch, or reload the image on specific module in the switch.

Reloading the Switch

To reload the switch, issue the **reload** command without any options. When you issue this command, you reboot the switch (see [Chapter 5, “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, if the identified module resides in slot 2:

```
switch# reload module 2
```

Reloading Switching Modules

Switching modules automatically download their images from the supervisor module, and do not need a force 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. For example, if 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...
```

Send documentation comments to mdsfeedback-doc@cisco.com

Preserving Module Configuration

To save the configuration, enter the **copy running-config startup-config** command from the EXEC mode prompt 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 6-3 displays various scenarios when module configurations are persevered or lost.

Table 6-3 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 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. | 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. |
| <p>Sample scenario:</p> <ol style="list-style-type: none"> 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. | <p>Sample response:</p> <ol style="list-style-type: none"> 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Purging Module Configuration

To delete the configuration in a specific module, issue the **purge module slot running-config** command from EXEC mode. Once this command is issued, the running configuration is cleared for the specified slot. This command will not work on supervisor modules or on any slot which 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, if you create an IP storage configuration with a an IPS module in slot 3 in Switch A. This module uses the 10.1.5.500 IP address. You decide to remove this IPS module and move it to Switch B, and you no longer need the 10.1.5.500 IP address. If you try to configure this unused 10.1.5.500 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 and then proceed with using this IP address.

Powering Off Switching Modules

By default, all switching modules are configured to be 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Identifying Module LEDs

Table 6-4 to Table 6-7 describe the LED location, type, and status for supervisor and switching modules used in Cisco MDS 9000 Family switches.

Table 6-4 *Module LEDs on a Cisco MDS 9200 Series Switch*

| Module | LED Type | Status | Description |
|---------------------------|----------|--------|---|
| Fixed switching module | Status | Green | <ul style="list-style-type: none"> All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) are reporting OK. Sufficient power is available for all modules |
| | | Orange | <ul style="list-style-type: none"> Any one of the chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) failed. Sufficient power is not available for all modules. Incompatible power supplies are installed. The redundant clock failed. |
| | | Red | <ul style="list-style-type: none"> The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. A temperature condition occurred. (A major threshold was exceeded during environmental monitoring.) |
| Optional switching module | System | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | | Orange | <ul style="list-style-type: none"> The module is booting or running diagnostics (normal initialization sequence). An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.) |
| | | Red | <ul style="list-style-type: none"> The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. An over temperature condition occurred. (A major threshold was exceeded during environmental monitoring.) |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 6-5 Supervisor Module LEDs on a Cisco MDS 9500 Series Switch

| LED | Status | Description |
|-----------------------|--------|--|
| Status | Green | All diagnostics pass. The module is online. |
| | Orange | <ul style="list-style-type: none"> The module is booting or running diagnostics (normal initialization sequence). The module is not online. An over temperature condition has occurred. (A minor threshold has been exceeded during environmental monitoring.) |
| | Red | <ul style="list-style-type: none"> The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. An over temperature condition has occurred. (A major threshold has been exceeded during environmental monitoring.) |
| System ¹ | Green | All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) are reporting OK. |
| | Orange | <ul style="list-style-type: none"> Any one of the environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) has failed. Incompatible power supplies are installed. 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. |

1. The System and Pwr Mgmt LEDs on a redundant supervisor module are synchronized to the active supervisor module.

Table 6-6 Ethernet Interface LEDs on a Cisco MDS 9200 Series Switch

| Module | LED Type | Status | Description |
|-------------------|----------|----------------|---|
| Ethernet (mgmt 0) | Activity | Flashing green | Traffic is passing through the interface. |
| | | Solid green | The link is functioning. |
| | | Off | The link is down. |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 6-7 Switching Module LEDs

| LED Type | Status | Description |
|----------|-----------------|---|
| Status | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | Orange | <ul style="list-style-type: none"> The module is booting or running diagnostics (normal initialization sequence). An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.) |
| | Red | <ul style="list-style-type: none"> The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. An over temperature condition occurred. (A major threshold was exceeded during environmental monitoring.) |
| Speed | On | 2 Gbps mode. |
| | Off | 1 Gbps mode. |
| Link | Solid green | Link is up. |
| | Flashing green | Link is up (beacon used to identify port). See the “Identifying the Beacon LEDs” section on page 9-12. |
| | Solid yellow | Disabled by software. |
| | Flashing yellow | Fault is detected. |
| | Off | Link is down. |

Default Supervisor Module Settings

[Table 6-8](#) lists the default settings for the supervisor module.

Table 6-8 Default Supervisor Module Settings

| Parameters | Default |
|----------------------------|---|
| Administrative connection | Serial connection. |
| Global switch information | <ul style="list-style-type: none"> 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 is set to 0.0.0.0. |



Managing System Hardware

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying Switch Hardware Inventory, page 7-2](#)
- [Displaying Power Usage Information, page 7-5](#)
- [Configuring Power Supplies, page 7-6](#)
- [Displaying Module Temperature, page 7-9](#)
- [Monitoring Fan Modules, page 7-10](#)
- [Monitoring Clock Modules, page 7-10](#)
- [Displaying Environment Information, page 7-11](#)

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Switch Hardware Inventory

Use the **show hardware** command to display switch hardware inventory details. See [Example 7-1](#).



Note

To display and configure modules, see [Chapter 6, “Managing Modules.”](#)

Example 7-1 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) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.

Software
  BIOS:          version 1.0.3
  loader:        version error [last 1.0(1)]
  kickstart:     version 1.1(1) [build 1.1(0.94)] [gdb]
  system:        version 1.1(1) [build 1.1(0.94)] [gdb]

  BIOS compile time:      11/18/02
  kickstart image file is: bootflash:/bootimage
  kickstart compile time: 2/12/2003 11:00:00
  system image file is:   isanimage
  system compile time:    2/12/2003 12:00:00

Hardware
  RAM 1027628 kB

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

  172.22.90.171 uptime is 0 days 3 hours 0 minute(s) 36 second(s)

  Last reset at 669882 usecs after Thu Feb 13 07:20:41 2003
  Reason: Reset Requested by CLI command reload
  System version: 1.0(1)

This supervisor carries Pentium processor with 1027628 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 ok
  Module type is "1/2 Gbps FC Module"
  1 submodules are present
  RAM size is 0 (kb)
  Model number is DS-X9016
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
H/W version is 0.0
Part Number is 73-8127-03
Part Revision is
Manufacture Date is Year 0 Week 0
Serial number is
CLEI code is

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 0.0
  Part Number is 73-7523-06
  Part Revision is
  Manufacture Date is Year 0 Week 0
  Serial number is
  CLEI code is

Module in slot 6 is empty

Module in slot 7 is empty

Module in slot 8 is ok
  Module type is "IP Storage Module"
  3 submodules are present
  RAM size is 0 (kb)
  Model number is DS-X9308-SMIP
  H/W version is 0.2
  Part Number is 73-8083-02
  Part Revision is 2
  Manufacture Date is Year 7 Week 2
  Serial number is JAB0702065h
  CLEI code is 0

Module in slot 9 is ok
  Module type is "1/2 Gbps FC Module"
  1 submodules are present
  RAM size is 0 (kb)
  Model number is DS-X9016
  H/W version is 0.0
  Part Number is 73-8127-03
  Part Revision is
  Manufacture Date is Year 0 Week 0
  Serial number is
  CLEI code is

-----
Chassis has 2 Slots for Power Supplies
-----

PS in slot A is ok
  Power supply type is "1153.32W 110v AC"
  Model number is DS-CAC-2500W
  H/W version is 1.0
  Part Number is 341-0061-01
  Part Revision is A0
  Manufacture Date is Year 6 Week 16
  Serial number is ART061600VA
  CLEI code is

PS in slot B is ok
  Power supply type is "1153.32W 110v AC"
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
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 20
Serial number is ART0620005N
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
```

Send documentation comments to mdsfeedback-doc@cisco.com

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 7-2](#).



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 7-2 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
-----
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Power Supplies

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**—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**—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 1200Watts when powered at 110 volts, and 2500Watts when powered at 220 volts.

To configure the power supply mode, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# confi g 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. |

Power Supply Guidelines



Note

Use the **show environment power** command to view the current power supply configuration.

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:
 - In **redundant** mode, the total power is the lesser of the two power supply capacities. For example, if you have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = not used

Current usage = 2000Watts

Current capacity = 2500Watts

Then the following three scenarios differ as specified (see [Table 7-1](#)):

- a. **Scenario 1:** If 1800Watts is added as power supply 2, then power supply 2 is shut down. Reason: 1800Watts is less than the usage of 2000Watts.
- b. **Scenario 2:** If 2200Watts is added as power supply 2, then the current capacity decreases to 2200Watts. Reason: 2200Watts is the lesser of the two power supplies.

Send documentation comments to mdsfeedback-doc@cisco.com

- c. **Scenario 3:** If 3000Watts is added as power supply 2, then the current capacity value remains at 2500Watts.

Reason: 2500Watts is the lesser of the two power supplies.

Table 7-1 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 2200Watts. |
| 3 | 2500 | 2000 | 3300 | 2500 | Capacity remains the same. |

1. W = Watts

- In **combined** mode, the total power is twice the lesser of the two power supply capacities.

For example, if you have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = not used

Current Usage = 2000Watts

Current capacity = 2500Watts

Then, the following three scenarios differ as specified (see [Table 7-2](#)):

- Scenario 1:** If 1800Watts is added as power supply 2, then the capacity increases to 3600Watts.
Reason: 3600Watts is twice the minimum (1800Watts).
- Scenario 2:** If 2200Watts is added as power supply 2, then the current capacity increases to 4400Watts.
Reason: 4400Watts is twice the minimum (2200Watts).
- Scenario 3:** If 3000Watts is added as power supply 2, then the current capacity increases to 5000Watts.
Reason: 5000Watts is twice the minimum (2500Watts).

Table 7-2 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

Send documentation comments to mdsfeedback-doc@cisco.com

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. Various configuration scenarios are displayed and summarized in [Table 7-3](#).

- a. **Scenario 1:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 1800Watts

Current Usage = 2000Watts

Current mode = **combined** mode (so current capacity is 3600Watts)

You decide to change the switch to **redundant** mode. Then power supply 2 is shut down.

Reason: 1800Watts is the lesser of the two power supplies and it is less than the system usage.

- b. **Scenario 2:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 2200Watts

Current Usage = 2000Watts

Current mode = **combined** mode (so current capacity is 4400Watts).

You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2200Watts.

Reason: 2200Watts is the lesser of the two power supplies.

- c. **Scenario 3:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 1800Watts

Current Usage = 3000Watts

Current mode = **combined** mode (so current capacity is 3600Watts).

You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2500Watts and the configuration is rejected.

Reason: 2500Watts is less than the system usage (3000Watts).

Table 7-3 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 | Existing configuration. |
| | 2500 | N/A | 2000 | 1800 | redundant | 2500 | Power supply 2 is shut down |
| 2 | 2500 | combined | 2000 | 2200 | N/A | 4400 | Existing configuration. |
| | 2500 | N/A | 2000 | 2200 | redundant | 2200 | The new capacity is changed. |
| 3 | 2500 | combined | 3000 | 1800 | N/A | 3600 | Existing configuration. |
| | 2500 | N/A | 3000 | 1800 | redundant | N/A | Rejected, so the mode reverts to combined mode. |

1. W = Watts

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Module Temperature

Use the **show environment temperature** command to display temperature sensors for each module (see [Example 7-3](#)).

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 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:
 - Syslog 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 as follows:
 - For sensors 1, 3, and 4 (outlet and onboard sensors):

Syslog 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, the module is shut down.

If the threshold is exceeded in a supervisor module with HA-standby or standby present, the supervisor module is shut down.

If the standby supervisor is not present, the entire switch is shut down.



Note

Switch shut down only happens after a two-minute interval. During this interval the software monitors the temperature every five (5) seconds and continuously sends syslog messages as configured. If the required action is not taken (for example, a new fan module is inserted to decrease temperature) and if the temperature does not come down, the system is shut down at the end of two minutes.

Example 7-3 Displays Temperature Information

```
switch# show environment temperature
```

| Module | Sensor | MajorThresh (Celsius) | MinorThres (Celsius) | CurTemp (Celsius) | Status |
|--------|--------|--------------------------|-------------------------|----------------------|--------|
| 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 |

Send documentation comments to mdsfeedback-doc@cisco.com

| | | | | | |
|---|--------|----|----|----|----|
| 7 | Outlet | 75 | 60 | 33 | ok |
| 7 | Intake | 65 | 50 | 30 | ok |
| 9 | Outlet | 75 | 60 | 34 | ok |
| 9 | Intake | 65 | 50 | 39 | ok |

Monitoring Fan Modules

Use the **show environment fan** command to display the fan status for each fan module. See [Example 7-4](#).

Example 7-4 Displays Chassis Fan Information

```
switch# show environment fan
```

| FAN | Model | Hw | Status |
|---------|--------------|-----|--------|
| Chassis | WS-9SLOT-FAN | 0.0 | ok |
| PS-1 | -- | -- | ok |
| PS-2 | -- | -- | ok |

The fan status is continuously monitored. In case of a fan module failure, the following action is taken:

- Syslog messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).



Caution

A fan failure could lead to temperature alarms if not corrected immediately.

Monitoring Clock Modules

Use the **show environment clock** command to display the clock status for the chassis. See [Example 7-5](#).

Example 7-5 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 |

Each switch has two clock modules for redundancy: Clock A (primary) and Clock B. The redundant clock module (Clock B) takes over if the primary clock module fails. If Clock A is available at startup, the switch uses Clock A, otherwise it uses Clock B.

If Clock A fails, the switch is reset and Clock B automatically takes over. Clock modules cannot be configured. If both modules fail, the switch shuts down. The probability of a clock failure is low given that the mean time between failures (MTBF) is 3660316 hours.

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Environment Information

Use the **show environment** command to display all environment-related switch information.

Example 7-6 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 (Celsius) | MinorThres (Celsius) | CurTemp (Celsius) | Status |
|--------|--------|--------------------------|-------------------------|----------------------|--------|
| 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 (Watts) | Power (Amp @42V) | Status |
|----|------------|------------------|---------------------|--------|
| 1 | PWR-950-AC | 919.38 | 21.89 | ok |
| 2 | | -- | -- | absent |

| Mod | Model | Power Requested (Watts) | Power Requested (Amp @42V) | Power Allocated (Watts) | Power Allocated (Amp @42V) | Status |
|-----|-----------------|-------------------------------|----------------------------------|-------------------------------|----------------------------------|------------|
| 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

```

Send documentation comments to mdsfeedback-doc@cisco.com



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 which allows identical Fibre Channel IDs (FCIDs) to be used simultaneously in different VSANs. 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 is provided (to another VSAN in the same physical VSAN) by a configured 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.

This chapter includes the following sections:

- [How VSANs Work, page 8-2](#)
- [VSANs Versus Zones, page 8-4](#)
- [Default and Isolated VSANs, page 8-6](#)
- [VSAN Membership, page 8-6](#)
- [VSAN Attributes, page 8-7](#)
- [Creating and Configuring VSANs, page 8-7](#)
- [Assigning VSAN Membership, page 8-8](#)
- [Deleting VSANs, page 8-9](#)
- [Viewing VSAN Configurations, page 8-10](#)
- [Default Settings, page 8-11](#)

Send documentation comments to mdsfeedback-doc@cisco.com

How VSANs Work


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.

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 (FCIDs) 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.

Figure 8-1 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. Within each VSAN, all members can talk to one another. Between VSANs no communication is possible.

Figure 8-1 Logical VSAN Segmentation




Send documentation comments to mdsfeedback-doc@cisco.com

Figure 8-2 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.

As displayed in both Figure 8-2 and Figure 8-2, the switch icons indicate that these features apply to any switch in the Cisco MDS 9000 family.

Figure 8-2 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 8-2 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

Send documentation comments to mdsfeedback-doc@cisco.com

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

VSANs Versus Zones

You can define multiple zones in a VSAN. Because two VSANs are equivalent to two nonconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. [Table 8-1](#) lists the differences between VSANs and zones.

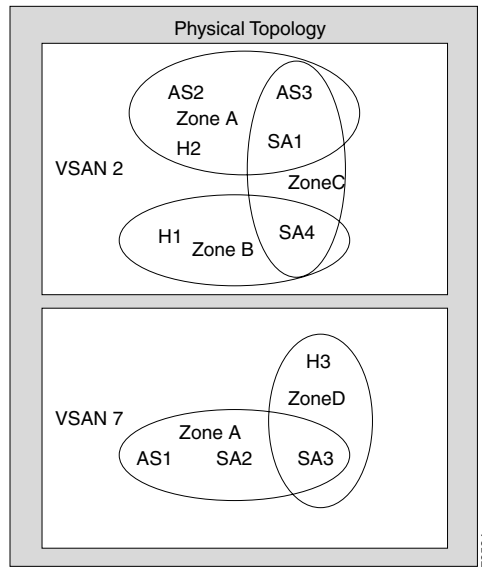
Table 8-1 VSAN and Zone Comparison

| VSANs | Zones |
|---|---|
| VSANs equal SANs with routing, naming, and zoning protocols. | These 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 may 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 8-3](#) 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 8-3 VSANS with Zoning



Send documentation comments to mdsfeedback-doc@cisco.com

Default and Isolated VSANs

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.

Default VSANs

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. If you do not need more than one VSAN for a switch, use this default VSAN as the implicit parameter during configuration. 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. It can be suspended.

Isolated VSANs

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.

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

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 change the VSAN membership by using the **vsan number interface type port/slot** command.

Trunking ports have an associated list of VSANs that are part of an allowed list (see [Chapter 10](#), “Configuring Trunking”).

Send documentation comments to mdsfeedback-doc@cisco.com

VSAN Attributes

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.

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.

Creating and Configuring VSANs

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

To create and configure 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 switch(config-vsan-db)# | 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 switch(config-vsan-db)# | Updates the VSAN with the assigned name (TechDoc). |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|---------|--|--|
| Step 5 | switch(config-vsan-db) # vsan 2 loadbalancing src-dst-id switch(config-vsan-db) # | Enables the load balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process. |
| Step 6 | switch(config-vsan-db) # no vsan 2 loadbalancing src-dst-id switch(config-vsan-db) # | Negates the command issued in the previous step and reverts to the default values of the load-balancing parameters. |
| Step 7 | switch(config-vsan-db) # vsan 2 loadbalancing src-dst-ox-id switch(config-vsan-db) # | Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default). |
| Step 8 | switch(config-vsan-db) # vsan 2 suspend switch(config-vsan-db) # | Suspends the selected VSAN. |
| Step 9 | switch(config-vsan-db) # no vsan 2 suspend vs.-config-vsan-db# | Negates the suspend command issued in the previous step. |
| Step 10 | switch(config-vsan-db) # end switch# | Returns you to EXEC mode. |

Assigning VSAN Membership

To assign VSAN membership, 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 switch(config-vsan-db) # | 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 switch(config-vsan-db) # | Assigns the membership of the fc1/8 interface to the specified VSAN (VSAN 2). |
| Step 5 | switch(config-vsan-db) # vsan 7 switch(config-vsan-db) # | Creates another VSAN with the specified ID (7) if that VSAN does not exist already. |
| Step 6 | switch(config-vsan-db) # vsan 7 interface fc1/8 switch(config-vsan-db) # | Updates the membership information of the interface to reflect the changed VSAN. |

Send documentation comments to mdsfeedback-doc@cisco.com

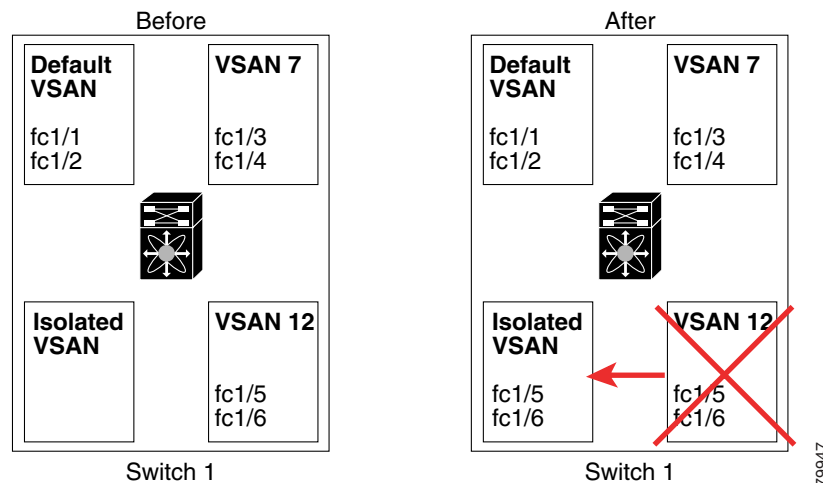
Deleting VSANs

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.

- VSAN attributes and port membership details are maintained by 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 8-4](#)).

Figure 8-4 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 10, “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.

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. |
| 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Viewing VSAN Configurations

Use the **show vsan** command to display information about configured VSANs (see Examples 8-1 to 8-6).

Example 8-1 *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 8-2 *Displays the VSAN Usage*

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

Example 8-3 *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
```

Example 8-4 *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 8-5 *Displays 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:
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
vsan 7 interfaces:
    fc1/8
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

Example 8-6 *Displays Membership Information for a Specified Interface*

```
switch # show vsan membership interface fc1/1
fc1/1
    vsan:1
    allowed list:1-4093
```

Default Settings

Table 8-2 lists the default settings for all configured VSANs.

Table 8-2 *Default VSAN Parameters*

| Parameters | Default |
|--------------------------|--|
| 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). |
| Port membership | Default VSAN (VSAN 1). |

Send documentation comments to mdsfeedback-doc@cisco.com



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, 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:

- [Configuring Fibre Channel Interfaces, page 9-2](#)
- [Default Settings, page 9-13](#)
- [Configuring the Management Interface, page 9-14](#)
- [Configuring VSAN Interfaces, page 9-15](#)
- [Displaying Interface Information, page 9-15](#)



Note

See [Chapter 3, “Initial Configuration”](#) and [Chapter 16, “Configuring IP Services,”](#) for more information on configuring mgmt0 interfaces.

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 3-15).

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Fibre Channel Interfaces

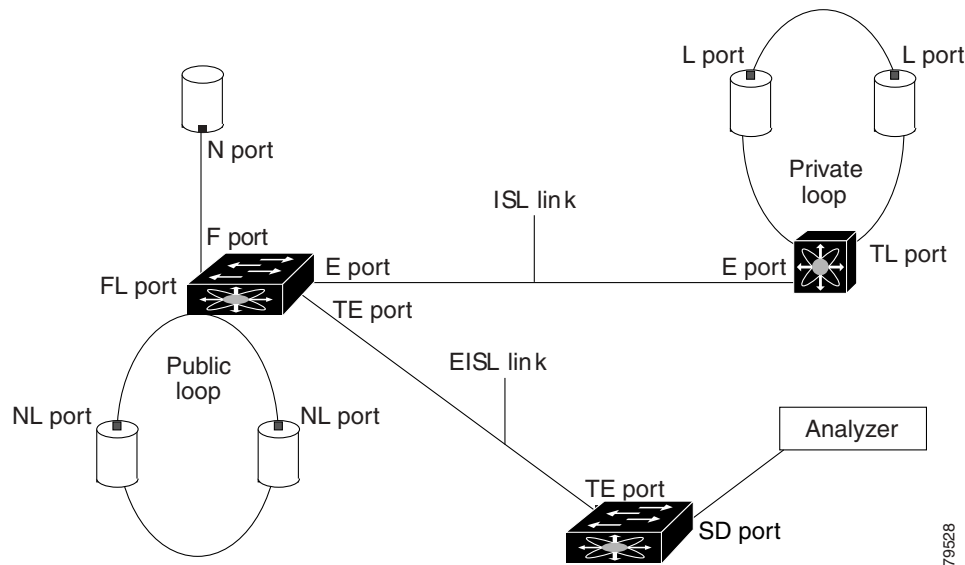
This section describes Fibre Channel interface characteristics, including (but are not limited to) modes, states, and speeds. It includes the following sections:

- [About Interface Modes, page 9-2](#)
- [About Interface States, page 9-5](#)
- [Configuring Fibre Channel Interfaces, page 9-8](#)
- [Configuring a Range of Interfaces, page 9-8](#)
- [Disabling Interfaces, page 9-8](#)
- [Configuring Interface Modes, page 9-9](#)
- [Configuring Administrative Speeds, page 9-9](#)
- [Configuring Interface Descriptions, page 9-10](#)
- [Configuring Buffer-to-Buffer Credits, page 9-10](#)
- [Configuring the Beacon Mode, page 9-11](#)

About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several modes: E port, F port, FL port, TL port, TE port, and SD port (see [Figure 9-1](#)). Besides these modes, each interface may be configured in auto or Fx port mode. These two modes determine the port type during interface initialization. A brief description of each interface mode follows.

Figure 9-1 Cisco MDS 9000 Family Switch Interface Modes



Note

Interfaces are created in VSAN 1 by default. See [Chapter 8, “Configuring and Managing VSANs.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com

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).

A brief description of each mode interface 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 11, “Configuring PortChannels”](#)).

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.

TL Port

In translatable loop port (TL port) mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices (NL ports). TL port mode is 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

See the [“Displaying TL Port Information” section on page 9-22](#). TL ports support class 2 and class 3 services.

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

Devices attached to TL ports are recommended to be configured in zones which have up to 64 zone members.

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 10, “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 through 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 23, “Monitoring Network Traffic Using SPAN”](#)).

Fx Port

Interfaces configured as Fx ports are allowed to 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 Cisco's 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 17-11](#) depicts a typical SAN extension over an IP network

When an FCIP peer is a SAN extender device that only support 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 17, “Configuring IP Storage”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Auto Mode

Interfaces configured as **auto** are allowed to 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 10, “Configuring Trunking”](#)). TL ports and SD ports are not determined during initialization and are administratively configured.

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 9-1](#).

Table 9-1 Administrative States

| Administrative State | Description |
|----------------------|--|
| Up | Enables an interface. |
| Down | Disables an interface. When an interface is administratively disabled (shutdown command), 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 9-2](#).

Table 9-2 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 9-3](#).

Send documentation comments to mdsfeedback-doc@cisco.com

Table 9-3 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 9-4 . |

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 9-4](#).

Table 9-4 Reason Codes for Nonoperational States

| Reason Code | Description | Applicable Modes |
|--------------------------------|---|------------------|
| Link failure or not connected | Physical layer link is not operational. | All |
| Fcot not present | The Fibre Channel optical transmitter 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 | Waiting 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: <ul style="list-style-type: none"> 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 configure the interface as shutdown followed by no shutdown . | |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 9-4 Reason Codes for Nonoperational States (continued)

| Reason Code | Description | Applicable Modes |
|---|--|-----------------------------|
| Isolation due to ELP failure | Port negotiation failed. | Only E ports and TE ports |
| Isolation due to ESC failure | 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 other side 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 32-port Switching Modules

The 32-port 1/2-Gbps switching module contains 8 port groups of 4 ports each. When configuring these modules the following 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 in the shutdown state.
- If any of the other three ports are configured in a no shutdown state, you cannot configure the first port as an E port. The other three ports continue to remain in a no shutdown state.
- Generally, the default port mode is auto. The auto option is not allowed in a 32-port switching module.

Send documentation comments to mdsfeedback-doc@cisco.com

- The default port mode for 32-port switching modules is Fx (Fx negotiates to F or FL).

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 | Configures the specified interface. |

When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

Configuring a Range of Interfaces

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 | Configures the range of specified interfaces. |
| | | Note In this command, provide a space before and after the comma. |

Disabling Interfaces

Interfaces on a port are shut down by default (unless you modified the initial configuration). To enable traffic flow, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 | Configures the specified interface. |
| Step 3 | switch(config-if)# no shutdown switch(config-if)# | Enables traffic flow to administratively allow traffic when the no prefix is used (provided the operational state is up). |
| | switch(config-if)# shutdown switch(config-if)# | Shuts down the interface and disables traffic flow (default). |

Send documentation comments to mdsfeedback-doc@cisco.com

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)# | Configures the specified interface. |
| 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. |

Configuring Administrative Speeds

By default, the administrative speed for an interface is automatically calculated by the switch. To configure the administrative speed of the interface, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config-if)# switchport speed 1000 switch(config-if)# | Configures the administrative speed of the interface to 1000 Mbps. The number indicates the speed in megabits per second (Mbps). You can set the speed to 1000 Mbps (for 1Gbps interfaces), 2000 Mbps (for 2 Gbps interfaces), or auto (default). |
| | switch(config-if)# switchport speed auto switch(config-if)# | Reconfigures the factory default (auto) administrative speed of the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Interface Descriptions

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)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport description cisco-HBA2 | Configures the description of the interface. The string may be up to 80 characters long. |
| | switch(config-if)# no switchport description | Clears the description of the interface. |

Configuring 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, since switches must not drop frames. Buffer Credits are negotiated on a per-hop basis.

The receive BB_credit value may be configured for each FC interface. In most cases, you don't need to modify the default configuration.



Note

The receive BB_credit values depend on the module type and the port mode:

16-port switching modules: 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.

32-port switching modules: The default value is 12 for the Fx, E, and TE modes. These values cannot be changed.

To configure buffer-to-buffer credits to 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)# | Configures the specified interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|--|---|
| Step 3 | switch(config-if)# switchport fcrxbbcrcedit default switch(config-if)# | 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 fcrxbbcrcedit 5 switch(config-if)# | 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 fcrxbbcrcedit 5 mode E switch(config-if)# | 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 fcrxbbcrcedit 5 mode Fx switch(config-if)# | Assigns this value if the port is operating in F or FL mode. The range to assign BB_credits is between 1 and 255. |

Configuring Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces by issuing the **switchport fcrxbuFSIZE** command. The default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure data field size 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)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport fcrxbuFSIZE 2000 switch(config-if)# | 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. |

Configuring the 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. The **beacon** command has no effect on the operation of the interface.

To disable beacon mode for a specified interface or range of interfaces, 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 beacon switch(config-if)# | Enables the beacon mode for the interface. |
| | switch(config-if)# no switchport beacon switch(config-if)# | Disables the beacon mode for the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

Identifying the Beacon LEDs

Figure 9-2 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 9-2 Cisco MDS 9000 Family Switch Interface Modes

| | | | |
|----------|---|----------|---|
| 1 | Status LED (see the “ Identifying Module LEDs ” section on page 6-8) | 3 | Link LEDs (see the “ Identifying Module LEDs ” section on page 6-8) and speed LEDs (explained in this section). |
| 2 | 1/2-Gbps Fibre Channel port group (see the “ Configuring 32-port Switching Modules ” section on page 9-7) | 4 | Asset tag (See the <i>Cisco MDS 9000 Family Hardware Installation Guide</i>). |

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—beacon mode is disabled
- On (flashing green)—the beacon mode is enabled. The LED flashes at one-second intervals.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Switchport Defaults

You can configure default values for various switch port attributes. If you configure the following attributes, they 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# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# no system default switchport shutdown switch(config-if)# | 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-if)# | 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-if)# | Configures the default setting for administrative trunk mode state of an interface as Auto. (The factory default setting is trunk mode On). |

Default Settings

Table 9-5 lists the default settings for Fibre Channel interface parameters.

Table 9-5 Default Fibre Channel Interface Parameters

| Parameters | Default |
|----------------------|--|
| Interface mode | Auto |
| Interface speed | Auto |
| Administrative state | Shutdown (unless changed during initial setup) |
| Trunk mode | On (unless changed during initial setup) |
| Trunk-allowed VSANs | 1 to 4093 |
| Interface VSAN | Default VSAN (1) |
| Beacon mode | Off |

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring the Management Interface

You can remotely configure the switch through the management interface (mgmt0). To configure a connection remotely, you must configure the IP parameters (IP address, subnet mask, and default gateway) from the CLI so that the switch is reachable.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask.

To configure the mgmt0 Ethernet interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface mgmt0 switch(config-if)# | Configures the management Ethernet interface on the switch to configure the management interface. |
| Step 3 | switch(config-if)# ip address 172.16.1.2 255.255.0 | Enters the IP address and IP subnet mask for the interface specified in Step 2. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |
| Step 5 | switch(config-if)# exit switch(config)# | Returns to configuration mode. |
| Step 6 | switch(config)# ip default-gateway 1.1.1.4 switch(config)# | Configures the default gateway IP address. |
| Step 7 | switch(config)# exit switch# | Returns to EXEC mode. |
| Step 8 | switch# copy running-config startup-config | Saves your configuration changes to the file system. Note This step is optional. If you wish to save your configuration, you can issue this command at any time. |

The management port (mgmt0) is autosensing and operates as full duplex mode and 100 Mbps speed. The speed and mode cannot be configured.



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.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring 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 nonexistent VSANs.

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 using the **interface VSAN** command. This is not done automatically.
- If you delete the VSAN, the attached interface is automatically deleted.

To create a VSAN interface, follow these steps:

| | Command | Purpose |
|--------|---|----------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 5 switch(config-if)# | Configures a VSAN with the ID 5. |

You can configure each interface only in one VSAN.

After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) features (see [Chapter 16, “Configuring IP Services”](#)).

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 9-1 to 9-9.

Example 9-1 Displays All Interfaces

```
switch# show interface
.
.
.
fc3/1 is trunking
  Hardware is Fibre Channel
  Port WWN is 20:81:00:05:30:00:12:5e
  Peer port WWN is 22:01:00:05:30:00:12:9e
  Admin port mode is E, trunk mode is auto
  Port mode is TE
  Port vsan is 2
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Receive Buffer Size is 2112
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1-15)
  Trunk vsans (up) (1-15)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  5 minutes input rate 40 bits/sec, 5 bytes/sec, 0 frames/sec
  5 minutes output rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

2161 frames input, 182556 bytes, 0 discards
    0 CRC, 0 unknown class
    0 too long, 0 too short
2164 frames output, 139904 bytes, 0 discards
1 input OLS, 1 LRR, 1 NOS, 0 loop inits
2 output OLS, 1 LRR, 1 NOS, 0 loop inits
.
.
.
fc9/9 is up
    Hardware is Fibre Channel
    Port WWN is 22:09:00:05:30:00:12:5e
    Admin port mode is auto, trunk mode is auto
    Port mode is FL, FCID is 0xef0100
    Port vsan is 1
    Speed is 1 Gbps
    Receive B2B Credit is 16
    Receive Buffer 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
        5 frames input, 560 bytes, 0 discards
            0 CRC, 0 unknown class
            0 too long, 0 too short
        4 frames output, 524 bytes, 0 discards
        0 input OLS, 0 LRR, 0 NOS, 2 loop inits
        2 output OLS, 0 LRR, 1 NOS, 1 loop inits
.
.
.
sup-fc0 is up
    Hardware is Fibre Channel
    Speed is 1 Gbps
    74994 packets input, 8076884 bytes
        0 multicast frames, 0 compressed
        0 input errors, 0 frame, 0 overrun 0 fifo
    74991 packets output, 7689168 bytes, 0 underruns
        0 output errors, 0 collisions, 0 fifo
        0 carrier errors

mgmt0 is up
    Hardware is FastEthernet
    Address is 0005.3000.2c5a
    Internet address is 172.22.90.38/24
    MTU 1500 bytes, Speed is 100 Mbps
    9319 packets input, 738784 bytes
        0 multicast frames, 0 compressed
        0 input errors, 0 frame, 0 overrun 0 fifo
    150 packets output, 34090 bytes, 0 underruns
        0 output errors, 0 collisions, 0 fifo
        0 carrier errors

vsan1 is up, line protocol is up
    WWPN is 10:00:00:05:30:00:12:63, FCID is 0xef001e
    Internet address is 10.10.11.10/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 2 is trunking
    Hardware is Fibre Channel
    Port WWN is 24:02:00:05:30:00:26:1e

```

Send documentation comments to mdsfeedback-doc@cisco.com

```

Admin port mode is E, trunk mode is on
Port mode is TE
Port vsan is 1
Speed is 4 Gbps
Trunk vsans (admin allowed and active) (1-5)
Trunk vsans (up) (1-5)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
5 minutes input rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
5 minutes output rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
  3534 frames input, 251672 bytes, 0 discards
    0 CRC, 0 unknown class
    0 too long, 0 too short
  3534 frames output, 176108 bytes, 0 discards
  9 input OLS, 8 LRR, 0 NOS, 0 loop inits
  13 output OLS, 11 LRR, 11 NOS, 0 loop inits
.
.
.

```

You can also specify arguments to display interface information.

Example 9-2 *Displays a Range of Interfaces*

```

switch# show interface fc2/5 , fc2/9
switch# show int fc2/5 - 10
fc2/5 is trunking
  Hardware is Fibre Channel
  Port WWN is 20:45:00:05:30:00:26:1e
  Peer port WWN is 21:85:00:05:30:00:25:9e
  Admin port mode is E, trunk mode is on
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Receive Buffer Size is 2112
  Encapsulation is normal
  Beacon is turned off
  Belongs to port-channel 2
  Trunk vsans (admin allowed and active) (1-5)
  Trunk vsans (up) (1-5)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  5 minutes input rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
  5 minutes output rate 8 bits/sec, 1 bytes/sec, 0 frames/sec
    1542 frames input, 105748 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    1542 frames output, 76656 bytes, 0 discards
    5 input OLS, 5 LRR, 0 NOS, 0 loop inits
    8 output OLS, 6 LRR, 6 NOS, 0 loop inits

fc2/9 is up
  Hardware is Fibre Channel
  Port WWN is 20:49:00:05:30:00:26:1e
  Peer port WWN is 21:89:00:05:30:00:25:9e
  Admin port mode is E, trunk mode is off
  Port mode is E, FCID is 0x7c0000
  Port vsan is 3
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Receive Buffer Size is 2112
  Encapsulation is normal
  Beacon is turned off

```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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
  624 frames input, 37736 bytes, 0 discards
    0 CRC, 0 unknown class
    0 too long, 0 too short
  625 frames output, 30248 bytes, 0 discards
  3 input OLS, 3 LRR, 2 NOS, 0 loop inits
  7 output OLS, 5 LRR, 5 NOS, 0 loop inits

```

Example 9-3 Displays a Specific Interface

```

switch# show interface fc2/9
fc2/9 is up
  Hardware is Fibre Channel
  Port WWN is 20:49:00:05:30:00:26:1e
  Peer port WWN is 21:89:00:05:30:00:25:9e
  Admin port mode is E, trunk mode is off
  Port mode is E, FCID is 0x7c0000
  Port vsan is 3
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Receive Buffer Size is 2112
  Encapsulation is normal
  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
    624 frames input, 37736 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    625 frames output, 30248 bytes, 0 discards
    3 input OLS, 3 LRR, 2 NOS, 0 loop inits
    7 output OLS, 5 LRR, 5 NOS, 0 loop inits

```

Example 9-4 Displays a VSAN Interface

```

switch# show int 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

```

Example 9-5 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
-----

```


Send documentation comments to mdsfeedback-doc@cisco.com

```
port-channel 1    --
port-channel 5    --
port-channel 6    --
```

Example 9-6 Displays Interface Information in a Brief Format

```
switch# show interface brief
```

```
-----
Interface  Vsan    Admin  Admin  Status      Oper  Oper  Port-channel
           Mode    Trunk  Mode
-----
fc2/5      1      E      on     trunking    TE    2     port-channel 2
fc2/6      1      E      on     trunking    TE    2     port-channel 2
fc2/7      1      E      on     down        --    --    --
fc2/8      1      auto   on     fcotAbsent  --    --    --
fc2/9      3      E      off    up          E     2     --
fc2/12     3      E      on     down        --    --    port-channel 4
fc3/14     1      SD     --     up          SD    1     --
fc9/1      1      auto   on     fcotAbsent  --    --    --
fc9/9      1      auto   auto   up          FL    1     --
-----

Interface      Status      Speed
-----
sup-fc0        up          1
-----

Interface      Status  IP Address      Speed      MTU
-----
mgmt0          up      172.22.90.38/24  100 Mbps   1500
-----

Interface      Status  IP Address      Speed      MTU
-----
vsan1          up      10.10.11.10/24   1 Gbps     1500
vsan2          up      10.10.12.10/24   1 Gbps     1500
-----

Interface      Vsan    Admin  Admin  Status      Oper  Oper
           Mode    Trunk  Mode
-----
port-channel 1  1      off    noOperMembers  --    --
port-channel 2  1      on     trunking      TE    4
port-channel 3  3      off    noOperMembers  --    --
-----
```

Example 9-7 Displays 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

    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
.
.
.
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

```

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

Interfaces 9/8 and 9/9 are not trunking ports and display class 2, 3, and F information as well.

Example 9-8 Displays Interface Counters in Brief Format

```
switch# show interface counters brief
```

| Interface | Input (rate is 5 min avg) | | Output (rate is 5 min avg) | |
|-----------|---------------------------|-----------------|----------------------------|-----------------|
| | Rate Mbits/s | Total Frames | Rate Mbits/s | Total 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 Mbits/s | Total Frames | Rate Mbits/s | Total Frames |
| port-channel 1 | 0 | 0 | 0 | 0 |
| port-channel 2 | 0 | 3946 | 0 | 3946 |

Example 9-9 Displays Transceiver Information

```
switch# show interface transceiver
```

```
.
.
.
fc9/6 fcot is present
  name is CISCO-AGILENT
  part number is QFBR-5796L
  revision is
  serial number is A00156980
  basic id fields (bytes 0-63)
    0x03 0x04 0x07 0x00 0x00 0x00 0x00 0x20
    0x40 0x0C 0x05 0x01 0x15 0x00 0x00 0x00
    0x1E 0x0F 0x00 0x00 0x43 0x49 0x53 0x43
    0x4F 0x2D 0x41 0x47 0x49 0x4C 0x45 0x4E
    0x54 0x20 0x20 0x20 0x00 0x00 0x30 0xD3
    0x51 0x46 0x42 0x52 0x2D 0x35 0x37 0x39
    0x36 0x4C 0x20 0x20 0x20 0x20 0x20 0x20
    0x20 0x20 0x20 0x20 0x00 0x00 0x00 0x86
  extended id fields (bytes 64-95)
    0x00 0x1A 0x00 0x00 0x41 0x30 0x30 0x31
    0x35 0x36 0x39 0x38 0x30 0x20 0x20 0x20
    0x20 0x20 0x20 0x20 0x30 0x32 0x30 0x38
    0x32 0x30 0x20 0x20 0x00 0x00 0x00 0x44
  vendor specific data (bytes 96-127)
    0x00 0x00 0x06 0x36 0x31 0x8D 0x23 0xB5
    0x8E 0xC2 0x13 0x9E 0xAC 0x57 0x47 0xB8
    0xAB 0x37 0x19 0x00 0x00 0x00 0x00 0x00
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

0x00 0x00 0x00 0x00 0x01 0x7D 0x67 0x74

fc9/7 fcot is present but not supported
name is IBM
part number is IBM42P21SNY
revision is AA20
serial number is 53P1487000WDN
basic id fields (bytes 0-63)
  0x03 0x00 0x07 0x00 0x00 0x00 0x00 0x20
  0x40 0x0C 0x05 0x01 0x15 0x00 0x00 0x00
  0x1E 0x0F 0x00 0x00 0x49 0x42 0x4D 0x20
  0x20 0x20 0x20 0x20 0x20 0x20 0x20 0x20
  0x20 0x20 0x20 0x20 0x00 0x08 0x00 0x5A
  0x49 0x42 0x4D 0x34 0x32 0x50 0x32 0x31
  0x53 0x4E 0x59 0x20 0x20 0x20 0x20 0x20
  0x41 0x41 0x32 0x30 0x00 0x00 0x00 0x07
extended id fields (bytes 64-95)
  0x00 0x1A 0x05 0x05 0x35 0x33 0x50 0x31
  0x34 0x38 0x37 0x30 0x30 0x30 0x57 0x44
  0x4E 0x20 0x20 0x20 0x30 0x32 0x30 0x35
  0x31 0x30 0x30 0x31 0x00 0x00 0x00 0x12
vendor specific data (bytes 96-127)
  0x49 0x42 0x4D 0x20 0x53 0x46 0x50 0x53
  0x20 0x41 0x52 0x45 0x20 0x43 0x4C 0x41
  0x53 0x53 0x20 0x31 0x20 0x4C 0x41 0x53
  0x45 0x52 0x20 0x53 0x41 0x46 0x45 0x20
.
.
.

```

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 since 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 9-9](#)).

The **show tlport** command displays the TL port interface configurations. This command provides a list of all TL ports configured on a box 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 Examples 9-10 to 9-13.

Example 9-10 *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.

Send documentation comments to mdsfeedback-doc@cisco.com

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 second lists the port WWNs, the third lists the node WWNs for each device, the fourth identifies the device as a SCSI initiator or target, and the last column is the real FC ID of the device.

Example 9-11 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 | 0xffffc42 |
| 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 9-12 Displays TL Port Information for Private Devices

```
switch# show tlport int fc1/16 pri
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 9-13 Displays TL Port Information for Proxied Devices

```
switch# show tlport int fc1/16 prox
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 | 0xffffc42 |
| 0x02 | 21:00:00:e0:8b:01:95:e7 | 20:00:00:e0:8b:01:95:e7 | Initiator | 0x420100 |

Send documentation comments to mdsfeedback-doc@cisco.com


TL Port Translation Guidelines

Table 9-6 lists the TL port translations supported in Cisco MDS 9000 Family switches:

Table 9-6 Supported TL Port Translation

| Translation from | Translation to | Example (see Figure 9-3) |
|----------------------------|-------------------------|-----------------------------|
| 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 9-3 TL Port Translation Support Examples



Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- To be proxied to the private loop, fabric devices must be in the same zone as private loop devices.
- 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.



Configuring Trunking


This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

- [About Trunking, page 10-1](#)
- [About Trunking Protocol, page 10-2](#)
- [Configuring Trunk Modes, page 10-3](#)
- [Configuring Trunk-Allowed VSAN List, page 10-4](#)
- [Trunking Configuration Guidelines, page 10-6](#)
- [Displaying Trunking Information, page 10-7](#)
- [Default Settings, page 10-8](#)

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 Extended ISL (EISL) frame format (see [Figure 10-1](#)).

Figure 10-1 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 (see the [“Configuring Trunk Modes”](#) section on page 10-3).
- 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 (see the [“Configuring Trunk-Allowed VSAN List”](#) section on page 10-4).
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port (see the [“About Trunking Protocol”](#) section on page 10-2).

Send documentation comments to mdsfeedback-doc@cisco.com

About Trunking Protocol

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

- Dynamic negotiation of operational trunk mode (see the “[Configuring Trunk Modes](#)” section on [page 10-3](#)).
- Selection of a common set of trunk-allowed VSANs (see the “[Configuring Trunk-Allowed VSAN List](#)” section on [page 10-4](#)).
- Detection of a VSAN mismatch across an ISL (see the “[Trunking Configuration Guidelines](#)” section on [page 10-6](#)).

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 will not be affected—the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated 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.



Tip

To avoid inconsistent configurations, shut all E ports before 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). |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring Trunk Modes

By default, the trunk mode is enabled in all Fibre Channel interfaces. However, the trunk mode configuration takes effect only in E-port mode. You can configure the 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 resulting trunking state of the link and the port modes at both ends (see [Table 10-1](#)).

Table 10-1 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 |



Note

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

To configure the 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 switch(config-if)# | Enables the trunk mode for the specified interface. |
| | switch(config-if)# switchport trunk mode off switch(config-if)# | Disables the trunk mode for the specified interface. |
| | switch(config-if)# switchport trunk mode auto switch(config-if)# | Configures the trunk mode for the specified interface. The auto option provides automatic sensing for the interface. |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)


Configuring Trunk-Allowed VSAN List

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 10-2](#), 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 10-2](#).


Figure 10-2 Default Allowed -Active VSAN Configuration



Send documentation comments to mdsfeedback-doc@cisco.com

You can configure a select set of VSANs (from the allowed-active list) to control access to those VSANs in a trunking ISL. Using Figure 10-2 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 10-3).

Figure 10-3 Operational and Allowed VSAN Configuration



In Figure 10-3, the operational allowed list of VSANs between switches is as follows:

- Switch 1 and switch 2 include VSAN 1 and VSAN 3.
- Switch 2 and switch 3 include VSAN 1 and VSAN 2.
- Switch 3 and switch 1 include VSAN 1, 2, and 5.

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

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 switch(config-if)# | Changes the allowed list for the specified VSANs. |
| | switch(config-if)# switchport trunk allowed vsan add 5 updated trunking membership switch(config-if)# | Expands the specified VSAN (5) to the new allowed list. |
| | switch(config-if)# no switchport trunk allowed vsan 2-4 switch(config-if)# | Deletes VSANs 2, 3, and 4. |
| | switch(config-if)# no switchport trunk allowed vsan add 5 switch(config-if)# | Deletes the expanded allowed list. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 VSANs merging (see [Figure 10-4](#)).

Figure 10-4 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 10-5](#)).

Figure 10-5 Third-Party Switch VSAN Mismatch



VSANs 2 and 3 get effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies (see the *Cisco MDS 9000 Family Fabric Manager User Guide*).

Send documentation comments to mdsfeedback-doc@cisco.com

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 10-1 to 10-3.

Example 10-1 *Displays a Trunked Fiber 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 10-2 *Displays Trunking Protocol*

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

Example 10-3 *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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Default Settings

Table 10-2 lists the default settings for trunking parameters.

Table 10-2 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 11

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. Specifically, a PortChannel has the following functionality:

- Provides a point-to-point connection over an 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.

Cisco MDS 9000 Family of switches support 128 PortChannels with 16 interfaces per PortChannel.

This chapter discusses the PortChannel feature provided in the switch. This chapter includes the following sections:

- [PortChannel Examples, page 11-2](#)
- [About PortChanneling and Trunking, page 11-3](#)
- [About Load Balancing, page 11-4](#)
- [Creating a PortChannel, page 11-5](#)
- [Deleting a PortChannel, page 11-6](#)
- [Adding Interfaces to a PortChannel, page 11-6](#)
- [Deleting Interfaces from a PortChannel, page 11-8](#)
- [Considerations for PortChannel Configurations, page 11-8](#)
- [Viewing PortChannel Information, page 11-9](#)
- [Default Settings, page 11-11](#)

Send documentation comments to mdsfeedback-doc@cisco.com

PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. [Figure 11-1](#) 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 11-1 PortChannel Flexibility

Send documentation comments to mdsfeedback-doc@cisco.com


About PortChanneling and Trunking

PortChanneling enables several links to be combined into one aggregated link.

Trunking enables an ISL to carry (trunk) multiple VSANs. Trunking can only be configured on 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 11-2](#)).

See [Chapter 10, “Configuring Trunking”](#) for information on trunked interfaces.

Figure 11-2 PortChanneling and Trunking



PortChanneling and trunking are used separately across an ISL:

- PortChanneling—Interfaces can be channeled between E ports over multiple ISLs or between TE ports over multiple EISLs.
- Trunking—Trunking, which permits carrying VSAN IDs between switches, can be done only between TE ports over EISLs.

See [Chapter 8, “Configuring and Managing VSANs.”](#)

Both PortChanneling and trunking can be used between TE ports over EISLs.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

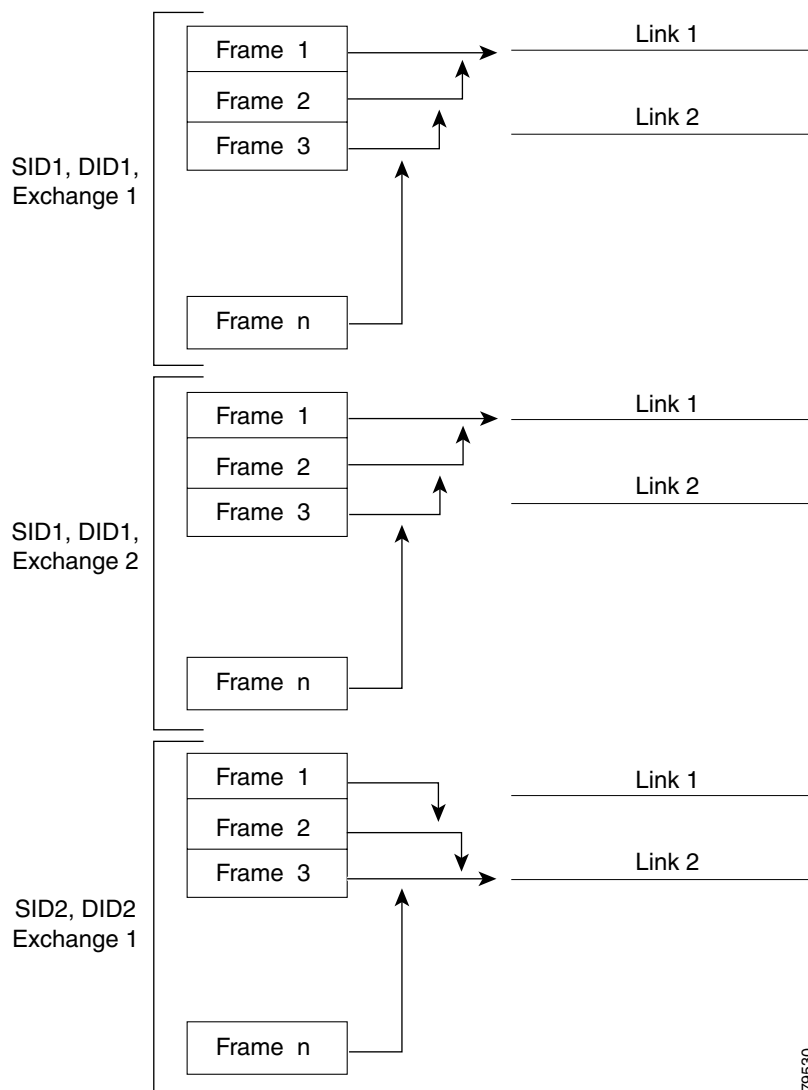
About Load Balancing

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 11-3 illustrates how source ID 1 (SID1) and destination ID1-based(DID1) 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 11-3 SID1 and DID1Based Load Balancing

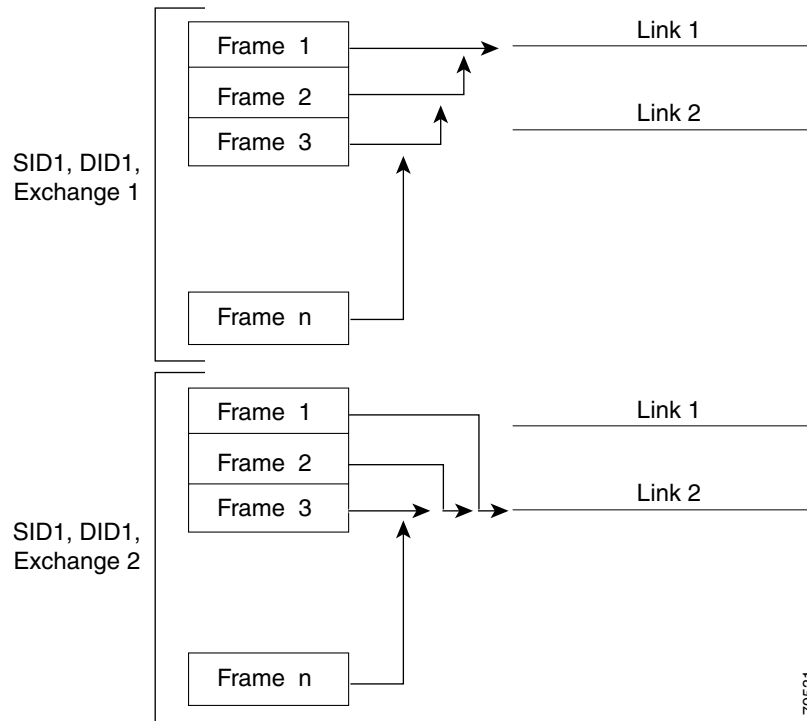


79530

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 11-4 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 11-4 SID1, DID1, and Exchange Based Load Balancing



For more information on configuring load balancing and in-order delivery features, see the “[VSAN Attributes](#)” section on page 8-7.

Creating a PortChannel

You can create PortChannels using the **interface port-channel** command. PortChannels are created with default values. You can change the default configuration just like any other physical interface.

To create a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface port-channel 1 switch(config-if)# | Configures the specified PortChannel (1). |



Note

All interfaces added to PortChannels are administratively shut down, and the PortChannel remains administratively up.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Deleting a PortChannel

To delete the PortChannel, you must explicitly issue the **no interface port-channel** command. When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. To avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. They continue to use the configured values of the physical port.

To delete a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | 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. |

Adding Interfaces to a PortChannel

You can add a physical interface (or a range of interfaces) to a nonexistent or an existing PortChannel and the PortChannel is automatically created. If the PortChannel does not exist, it is created. The compatible parameters on the configuration are mapped to the PortChannel.

To add a port (or a range of ports) to a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/15 switch(config-if)# | Configures the specified port interface (fc1/15). |
| | 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 |
| 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 switch(config-if)# | 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. |
| | 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 switch(config-if)# | 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Forcing an Interface Addition

You can specify a **force** option to force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel and the port is shut down.



Note

When PortChannels are created automatically, the **force** option cannot be used.

To force the addition of a port to a PortChannel, 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 port 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 a physical Fibre Channel port 1/1 addition to channel group 1. The E port is shut down. |

Compatibility Check

A compatibility check ensures that the same configuration values are used in all physical ports in the channel. For example, to enable trunk mode, all operational ports in the configuration must be configured in the trunk mode or in the nontrunking mode. Otherwise, they cannot become part of a PortChannel. A port cannot be operational if it is incompatible with the PortChannel. If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces.

Suspended State

An interface enters the suspended state if its operational values are incompatible with the PortChannel. A compatibility check on operational parameters is done when one of the following events occurs:

- A port becomes operational in a PortChannel.
- An operational parameter changes for a port in a PortChannel.

The software performs a compatibility check on the operational parameters and places the interface in an operational or suspended state based on the result (see the [“Reason Codes”](#) section on page 9-5).

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Deleting Interfaces from a PortChannel

To delete a physical interface (or a range of physical interfaces), you must explicitly issue the **no channel-group** command at the physical interface level. 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.



Note

When an interface is deleted, it is shut down but the physical configuration is retained. The inherited PortChannel configuration information is not deleted.

To delete a physical interface (or a range of physical interfaces), 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. |

Considerations for PortChannel Configurations

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 two switches. PortChannels require point-to-point connections.

Error Detection

If you misconfigure PortChannels, you may receive the `Error disabled - Possible port channel misconfiguration` message. If you receive this message, the PortChannel's physical links are disabled since 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.
- Links in a PortChannel must not 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 re-enable the links. Issue the **show interface** command for that interface to verify that the PortChannel is functioning as required.

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

Send documentation comments to mdsfeedback-doc@cisco.com

Viewing PortChannel Information

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.

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 interface (FOP), which is the primary operational interface selected in the PortChannel. See Examples 11-1 to 11-6.

Example 11-1 PortChannel Summary

```
switch# show port-channel summary
```

| Interface | Total Ports | Oper Ports | First Oper Port |
|----------------|-------------|------------|-----------------|
| port-channel 1 | 2 | 2 | fc2/3 |
| port-channel 2 | 2 | 2 | fc2/5 |
| port-channel 3 | 2 | 2 | fc2/10 |
| . | | | |
| . | | | |
| . | | | |

Example 11-2 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
```

Example 11-3 PortChannel Database

```
switch# show port-channel database
port-channel 1
  Administrative channel mode is on
  Operational channel mode is on
  Last membership update succeeded
  First operational port is fc2/3
  2 ports in total, 2 ports up
  Ports:  fc2/3    [up]
          fc2/4    [up]
port-channel 2
  Administrative channel mode is on
  Operational channel mode is on
  Last membership update succeeded
  First operational port is fc2/5
  2 ports in total, 2 ports up
  Ports:  fc2/5    [up]
          fc2/6    [up]
.
.
.
```

Send documentation comments to mdsfeedback-doc@cisco.com

The **show port-channel consistency** command has two options—without detail and **detail**.

Example 11-4 Command Without Details

```
switch# show port-channel consistency
sup database:
=====
totally 7 port-channels
port-channel 1:
    2 ports, first operational port is fc2/3
    fc2/3    [up]
    fc2/4    [up]
port-channel 2:
    2 ports, first operational port is fc2/5
    fc2/5    [up]
    fc2/6    [up]
.
.
.
```

Example 11-5 Command With Details

```
switch# show port-channel consistency detail
Authoritative port-channel database:
=====
totally 7 port-channels
port-channel 1:
    2 ports, first operational port is fc2/3
    fc2/3    [up]
    fc2/4    [up]
port-channel 2:
    2 ports, first operational port is fc2/5
    fc2/5    [up]
    fc2/6    [up]
.
.
.
=====
database 1: from module 5
=====
totally 7 port-channels
port-channel 1:
    2 ports, first operational port is fc2/3
    fc2/3    [up]
    fc2/4    [up]
port-channel 2:
    2 ports, first operational port is fc2/5
    fc2/5    [up]
    fc2/6    [up]
.
.
.
=====
database 3: from module 2
=====
totally 7 port-channels
port-channel 1:
    2 ports, first operational port is fc2/3
    fc2/3    [up]
    fc2/4    [up]
port-channel 2:
    2 ports, first operational port is fc2/5
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
fc2/5    [up]
fc2/6    [up]
.
.
.
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

Example 11-6 PortChannel Usage

```
switch# show port-channel usage
Totally 7 port-channel numbers used
=====
Used   :    1-7
Unused:    8-128
```

Default Settings

[Table 11-1](#) lists the default settings for PortChannels.

Table 11-1 Default PortChannel Parameters

| Parameters | Default |
|--------------------|-----------------------------|
| PortChannels | FSPF is enabled by default. |
| Create PortChannel | Administratively up. |
| Default mode | Auto. |

Send documentation comments to mdsfeedback-doc@cisco.com



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. This chapter defines various zoning concepts and provides details on zone set and management features in the switch and includes the following sections:

- [Zoning Features, page 12-2](#)
- [Zoning Example, page 12-3](#)
- [Configuring a Zone, page 12-4](#)
- [Configuring Aliases, page 12-4](#)
- [Zone Enforcement, page 12-5](#)
- [Zone Sets, page 12-5](#)
- [A Default Zone, page 12-9](#)
- [Recovering from Link Isolation, page 12-10](#)
- [Distributing Zone Sets, page 12-11](#)
- [Copying Zone Sets, page 12-11](#)
- [Clearing Zone Sets, page 12-11](#)
- [Viewing Zone Information, page 12-12](#)
- [Default Settings, page 12-15](#)
- [Zone Implementation, page 12-16](#)

Table 8-1 on page 8-4 lists the differences between zones and VSANs.



Note

For a comprehensive summary of zone implementation, refer to the “[Zone Implementation](#)” section on page 12-16.

Send documentation comments to mdsfeedback-doc@cisco.com

Zoning Features

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 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.
- 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 the option (**zoneset distribute full vsan** command) 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 configured without interrupting traffic on unaffected ports or devices.
- Zone membership criteria is based 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.
- 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Zoning Example


Figure 12-1 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 12-1 Fabric with Two Zones



Of course, there are other ways to partition this fabric into zones. Figure 12-2 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 12-2 Fabric with Three Zones



[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring a Zone

A zone can be configured using one of the following types to assign members:

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- Fabric port WWN—The WWN of the fabric port name in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID in 0xhhhhhh format (for example, 0xce00d1).
- FC alias—The alias name is in alphabetic characters (for example, Payroll) and denotes a port ID or WWN. The alias can also include multiple members.

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 Zone 1 for the VSAN called vsan3. |
| 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 | Configures a member for the specified zone (Zone1) based on the type (pWWN, fabric pWWN, FC ID, or FC alias) and value specified. |
| Tip | Use a relevant display command (for example, show interface or show flogi database) to obtain the required value in hex format. | |

Configuring Aliases

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), or pWWN values.

To create an alias using the **fcalias** command, 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). |
| Step 3 | switch-config-fcalias# member fcid 0x222222 switch-config-fcalias# | Configures alias members based on the specified FC ID type and value (0x222222). |
| | switch-config-fcalias# member pwwn 10:00:00:23:45:67:89:ab switch-config-fcalias# | Configures alias members based on the specified port WWN type and value (pWWN 10:00:00:23:45:67:89:ab). |
| | switch-config-fcalias# member fwwn 10:01:10:01:10:ab:cd:ef switch-config-fcalias# | Configures alias members based on the specified fWWN type and value (fWWN 10:01:10:01:10:ab:cd:ef). |

Send documentation comments to mdsfeedback-doc@cisco.com



Note

Multiple members can be specified on multiple lines.

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.



Note


Hard zoning enforces zoning restrictions on every frame, and it prevents unauthorized access at all times.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

Zone Sets

In [Figure 12-3](#), two separate sets are created, each with its own membership hierarchy and zone members.

Figure 12-3 Hierarchy of Zone Sets, Zones, and Zone Members



79537

Send documentation comments to mdsfeedback-doc@cisco.com

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 both at once).

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. Note To activate a zone set, you must first create the zone and a zone set. |
| Step 3 | switch-config-zoneset# member Zone1 switch-config-zoneset# | Adds Zone1 as a member of the specified zone set (Zoneset1). Note If the specified zone name was not previously configured, this command will return the <code>zone not present</code> error message. |
| 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 newly created zone (InlineZone1). Tip Execute this step only if you need to add a member to a zone from a zone set prompt. Note Multiple members can be specified on multiple lines. |

Zone sets are configured with the names of the member zones. If the zone set is in a configured VSAN, the VSAN is also specified.

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. You can activate a zone set using the **zoneset activate name** command.
- The administrator can modify the full zone set even if a zone set with the same name is active. The changes do not take effect until the zone set is activated with the **zoneset activate name** command.

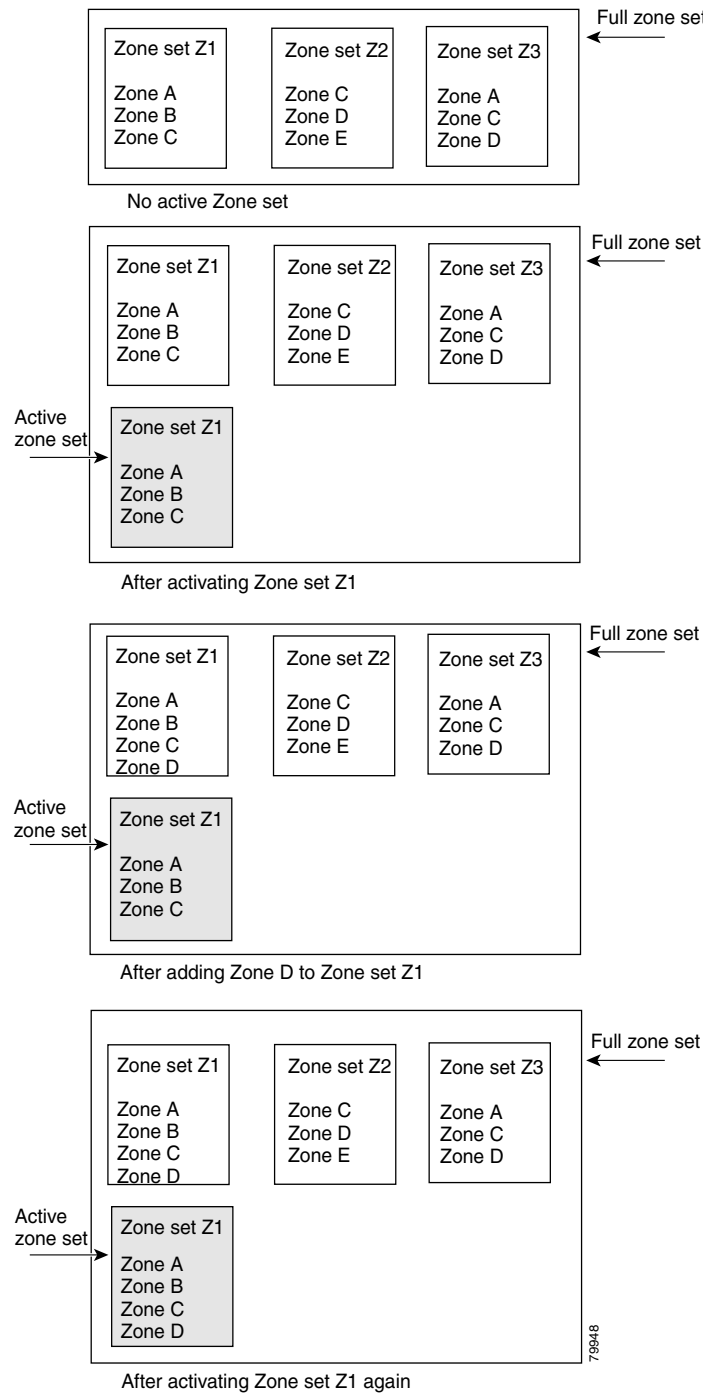
Send documentation comments to mdsfeedback-doc@cisco.com

- 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. 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.
- 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.

Figure 12-4 shows a zone being added to an activated zone set.

Send documentation comments to mdsfeedback-doc@cisco.com

Figure 12-4 Active and Full Zone Sets



To activate a zone set, follow these steps:

| | Command | Purpose |
|--------|-------------------------|----------------------------|
| Step 1 | switch# config t | Enters configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|--|------------------------------------|
| Step 2 | switch(config)# zoneset activate name Zoneset1 vsan 3 switch(config)# | Activates the specified zone set. |
| | switch(config)# no zoneset activate name Zoneset1 vsan 3 switch(config)# | Deactivates the specified zone set |

**Note**

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

A 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 a 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 be permitted or denied to members of the default zone. This information is not distributed to all switches; it must be performed for 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. The default zone members are explicitly listed only when the default policy is configured as **permit**. 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.

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 1 switch(config)# | Permits traffic flow to default zone members. |
| | switch(config)# no zone default-zone permit vsan 1 switch(config)# | Denies traffic flow to default zone members and reverts to factory default. |

**Note**

The default settings for default zone configurations can be changed.

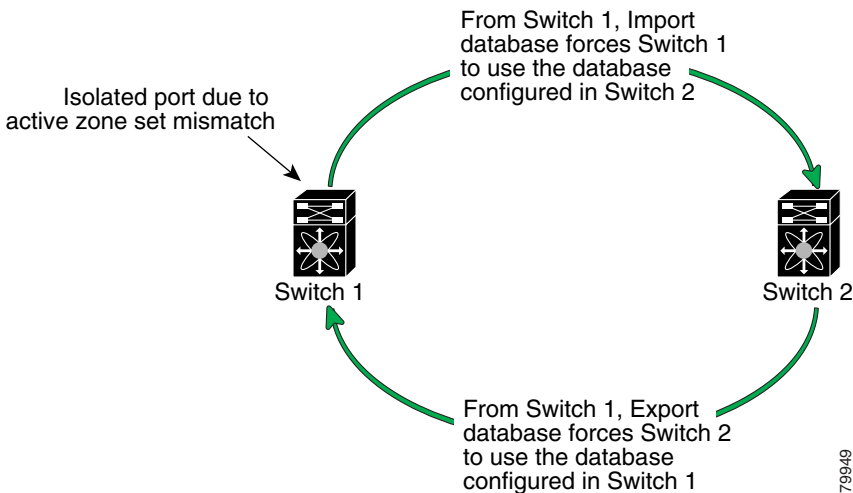
[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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 zoneset 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 zoneset database and replace the current active zoneset (see [Figure 12-5](#)).
- Export the current database to the neighboring switch (see [Figure 12-5](#)).
- Manually resolve the conflict by editing the full zone set, bringing up the link, and then activating the corrected zone set.

Figure 12-5 Importing and Exporting the Database




Tip

The **import** and **export** commands should be issued from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

To import the zone database from an adjacent switch, follow these steps:

| Step 1 | Command | Purpose |
|--------|--|--|
| | switch# zone merge interface fc1/3 import vsan2 | Imports the zone database from the adjacent switch connected through the VSAN 2 interface. |
| | switch# zone merge interface fc2/8 export vsan5 | Exports the zone database from the adjacent switch connected through the VSAN 5 interface. |


Note

You can also issue **zone merge interface** commands for a range of VSANs.

Send documentation comments to mdsfeedback-doc@cisco.com

Distributing Zone Sets

When a zone set is activated, by default, only the active zone set is sent to other switches in the Fabric. This command enables sending a full zone set along with the active zone set. It takes effect while sending merge request frame to the adjacent switch.

To distribute zone sets, 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. |

Copying Zone Sets

You can copy an active zone set to a full zone set using the **zone copy active-zoneset full-zoneset** command. You can not make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated. This command does not distribute zone sets.



Note

Since you can not edit an active zone set, this command is helpful in changing a copy of an existing zone set. You can make a copy and then edit it without altering the existing active zone set.

The **zone copy** command is used to copy active zone sets to the full zone set.

To distribute zone sets, follow this step:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# zone copy active-zoneset full-zoneset Please enter yes to proceed.(y/n) [n]? y | Makes a copy of the active zone set in the full zone set. |

Clearing Zone Sets



Note

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

To clear the zone server database, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# clear zone database vsan2 | Clears all configured information in the zone server for the specified VSAN. |



Note

After issuing a **clear 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Viewing Zone Information

You can view any zone information for any configured interface by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, alias, or even a keyword like **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 12-1 to 12-12.

Example 12-1 *Displays Zone Information for All VSANs*

```
switch(config)# 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
```

Example 12-2 *Displays Zone Information for a Specific VSAN.*

```
switch(config)# 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 12-3 *Display 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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 12-4 Display 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 12-5 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 12-6 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Example 12-7 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 12-8 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 12-9 Displays Active Zonesets

```
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 12-10 Displays Brief Descriptions of Zone Sets

```
switch# show zoneset brief
zoneset name ZoneSet1 vsan 1
  zone zone1
  zone zone2
```


Send documentation comments to mdsfeedback-doc@cisco.com

Example 12-11 Displays Active Zones

```
switch# show zone active
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 12-12 Displays Zone Status

```
switch# show zone status
VSAN: 1 default-zone: deny distribute: full Interop: Off
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
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
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
```

Default Settings

Table 12-1 lists the default settings for zone parameters.

Table 12-1 Default Zone Parameters

| Parameters | Default |
|--------------------------|--|
| Default zone policy | Denied to all members. |
| Full zone set distribute | The full zone set(s) is not distributed. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 changed (more secure).
- 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 zonesets 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

You can additionally configure the following zone features if required:

- Propagate full zone sets to all switches on a per VSAN basis using the **zoneset distribute full-database vsan** command.
- Change the default policy for unzoned members using the **zone default permit vsan** command.
- Inter-operate with other vendors by configuring a VSAN in the interop mode using the **vsan 1 interop** command.
 - 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 using the **export** or **import** commands



Managing FLOGI, Name Server, and RSCN Databases

This chapter describes the fabric login database, the name server features, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family. It includes the following sections:

- [Displaying FLOGI Details, page 13-1](#)
- [Configuring the Name Server Proxy Feature, page 13-3](#)
- [Displaying RSCN Information, page 13-6](#)

Displaying FLOGI Details

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. See Examples [13-1](#) to [13-4](#).

Example 13-1 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 13-2 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 |

Send documentation comments to mdsfeedback-doc@cisco.com

```

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    20:00:00:20:37:46:78:97

```

Total number of flogi = 10.

Example 13-3 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 13-4 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 “Allocating Flat FC IDs” section on page 24-18 and the “Enabling Loop Monitoring” section on page 24-18.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring the Name Server Proxy Feature

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 which 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.

All name server registration requests come from the same port whose parameter is registered or changed. If it doesn't, 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 switch(config)# | Configures a proxy port for the specified VSAN. |

Displaying 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.

Use the **show fcns** command to display the name server database and statistical information for a specified VSAN or for all VSANs (see Examples 13-5 to 13-9).

Example 13-5 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 | (Andiamo) | 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 | (Andiamo) | ipfc |
| 0x020200 | N | 21:01:00:e0:8b:22:99:36 | (Company A) | scsi-fcp |

Send documentation comments to mdsfeedback-doc@cisco.com

Example 13-6 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
0x030300      N      10:00:00:4a:c9:28:c7:01
0x030400      N      10:00:00:59:c9:28:c7:01
0xec0001      NL     21:00:00:20:37:a6:be:14 (Seagate)         scsi-fcp
0xec0100      N      10:00:00:05:30:00:26:23 (Cisco)           ipfc
0xec0200      N      10:00:00:5a:c9:28:c7:01

Total number of entries = 8
```

Example 13-7 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 = 8
```

Send documentation comments to mdsfeedback-doc@cisco.com

Example 13-8 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
```

Example 13-9 Displays the Internal Name Server Information for the Specified VSAN

```
switch# show fcns internal info vsan 8
Local Domain: 13
Remote Domain:
Info for 21
updating_db = 0
Requests sent to the switch with response pending:
Indexed objects details:
port_id index::
size:10240 incr_factor:512 slots_free:10239
portwn index::
size:10240 incr_factor:512 slots_free:10239
nodewwn index::
size:10240 incr_factor:512 slots_free:10239
ip_addr index::
size:10240 incr_factor:512 slots_free:10240
```

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying RSCN Information

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
- Or any other similar event that affects the operation of the host

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.

Use the **show rscn** command to display RSCN information (see Examples 13-10 and 13-11).

Example 13-10 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 cannot be configured, it is only populated if one or more hosts send SCR frames to register for RSCN information. If the **show rscn scr-table** command does not return any entries, no host is interested in receiving RSCN information.

Example 13-11 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
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
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
```

Sending RSCNs

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

For example, you have two disks (D1, D2) and a host (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:

- If the **multi-pid** option is disabled on switch 1, then two RSCNs is generated to Host H—one for the disk D1 and another for disk D2.
- If the **multi-pid** format is enabled on switch 1, then a single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

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 the specified VSAN. |



Note

Some Nx ports may not understand multi-pid RSCN payloads. If so, you must disable the **multi-pid** RSCN option.

Clearing RSCN Statistics

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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
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
```

This command is used for debugging purposes. When you 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.



Configuring System Security and AAA Services

Security can be independently configured for each of the following management paths:

- Command-line interface (CLI)—You can access the CLI using one of three connection options:
 - Console (serial connection)
 - Telnet
 - Secure Shell Protocol (SSH)
- Simple Network Management Protocol (SNMP)—The SNMP agent supports security features FOR versions 1, 2c, and 3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).



Note

Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for information on this management tool.

This chapter includes the following sections:

- [Management Security Features, page 14-2](#)
- [Authentication and Authorization Process, page 14-4](#)
- [Configuring CLI Authentication Methods, page 14-5](#)
- [Configuring Role-Based CLI Authorization, page 14-7](#)
- [Configuring CLI User Profiles, page 14-9](#)
- [Configuring CLI Accounting Parameters, page 14-11](#)
- [Recovering Administrator Password, page 14-13](#)
- [Configuring RADIUS Authentication, page 14-14](#)
- [Configuring SSH Services, page 14-18](#)
- [SNMP Security, page 14-20](#)
- [Default Security Settings, page 14-26](#)

Send documentation comments to mdsfeedback-doc@cisco.com

Management Security Features

Table 14-1 shows the security features of the Cisco MDS 9000 Family switches.

Table 14-1 Management Security Features

| Security Features | CLI (Console or Telnet/SSH Access) | SNMP (v1, v2c, and v3 access) |
|---|--|---|
| User authentication | Local and RADIUS | Local only |
| Role-based authorization | Local and RADIUS | Local only |
| Accounting | Local and RADIUS | Local and RADIUS (only logs configuration commands) |
| Encryption management access | SSH only (not applicable for console or Telnet access) | SNMPv3 |
| Anti-replay attack and prevention of man-in-middle attack | | |



Note

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

User Authentication

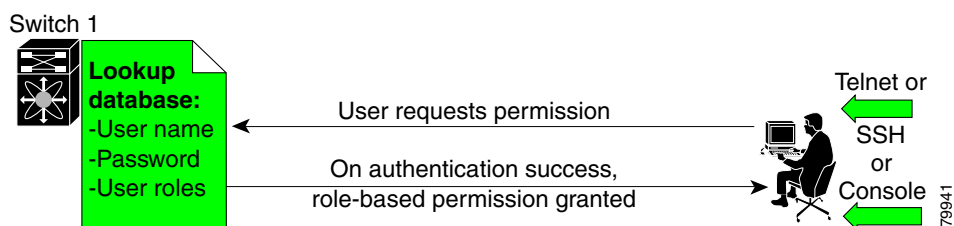
Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers).

For each management path (console or Telnet and SSH), you can enable only one of three options—local, RADIUS, or none. The option can be different for each path.

Local Authentication

The system maintains the user name and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored information (see Figure 14-1).

Figure 14-1 Local Authentication



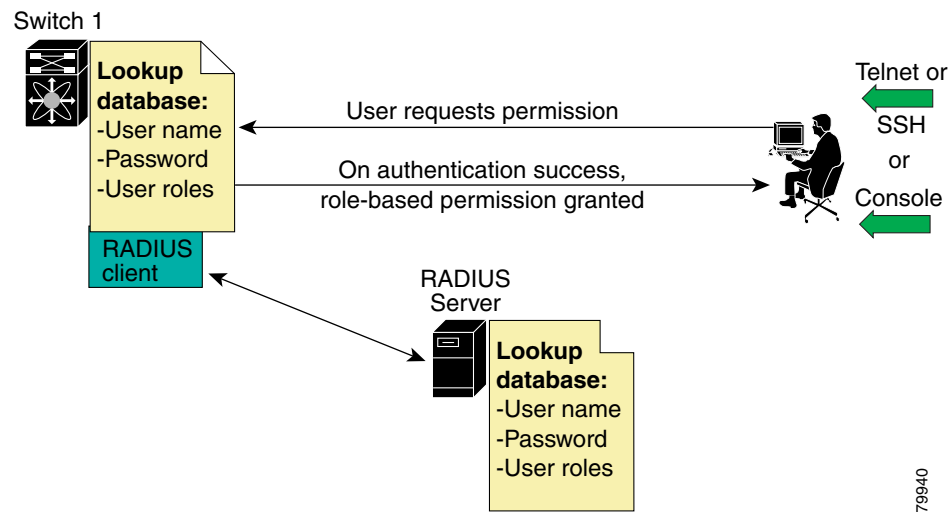
Send documentation comments to mdsfeedback-doc@cisco.com

RADIUS Authentication

Cisco MDS 9000 Family switches provide remote authentication through RADIUS servers. You can also configure multiple RADIUS servers, and each server is tried in the order specified.

RADIUS protocols support one-time password (OTP) schemes that all switches can make use of for authentication purposes (see [Figure 14-2](#)).

Figure 14-2 RADIUS Authentication



79940

Role-Based Authorization

By default, two roles exist in all 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.

The two default roles cannot be changed or deleted. Vendor-specific attributes (VSAs) contain the user profile information used by the switch. To use this option, configure the VSAs on the RADIUS servers.

Accounting

Accounting refers to the log that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally and remotely (using RADIUS).

Send documentation comments to mdsfeedback-doc@cisco.com

Authentication and Authorization Process

The following steps explain the authorization and authentication process. [Figure 14-3](#) shows a flow chart of the process.


- Step 1** The switch software receives your user ID and password through a console or Telnet (or SSH) application.
- Step 2** The remote server is contacted if remote authentication is enabled, or else local authentication is performed.



Note If remote authentication is enabled but none of the servers are available (network failure), local authentication is performed.

- Step 3** If authentication is successful, you are given access to the switch with appropriate permissions based on the roles to which you belong. These roles can be configured locally or can be sent by the remote server during the authentication process.
- Step 4** If remote authentication is rejected, you are denied access and an appropriate message is issued.

Figure 14-3 Switch Authorization and Authentication Flow



79951

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring CLI Authentication Methods

You can configure remote and local authentication for Telnet, SSH, or console access. These commands are restricted to privileged users as determined by your administrator.

Setting AAA Authentication

You can individually set authentication options for console or Telnet (and SSH) access using the **aaa authentication login** command. Local authentication is always disabled by default (see the “Authentication and Authorization Process” section on page 14-4).

To configure the authentication option, follow these steps:

| | Command | Purpose |
|--|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# aaa authentication login radius telnet switch(config)# | Enables Telnet authentication (and SSH) to use RADIUS. |
| | switch(config)# aaa authentication login radius console switch(config)# | Enables console authentication to use RADIUS. |
| | switch(config)# aaa authentication login local telnet | Enables only local authentication for Telnet (and SSH) access. |
| Note This command applies to both Telnet and SSH. | | The local option disables other authentication methods and configures local authentication to be used exclusively. |
| | switch(config)# aaa authentication login none console | Disables authentication for console access. User name authentication is still done. |

Enabling or Disabling Telnet Access

You can use the **telnet server enable** command to enable Telnet access to the switch. By default, this service is enabled.

To enable or disable Telnet access to the switch, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# telnet server enable switch(config)# | Turns on Telnet access to the switch. |
| | switch(config)# no telnet server enable switch(config)# | Turns off Telnet access to the switch. |

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying CLI Authentication Commands

The **show authentication** command displays the configured authentication methods. See [Example 14-1](#).

Example 14-1 Displays Authentication Information

```
switch# show authentication
authentication method:none
    console:not enabled
    telnet/ssh:not enabled
authentication method:radius
    console:not enabled
    telnet/ssh:not enabled
authentication method:local
    console:enabled
    telnet/ssh:enabled
```

The **show telnet server** command displays the state of the Telnet access configuration. See [Example 14-2](#).

Example 14-2 Displays Telnet Server Details

```
switch# show telnet server
telnet service enabled
```


Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Role-Based CLI 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.

To configure a new 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 | Includes all users in the Tech. Docs. group. The description is restricted to one line and can contain spaces. |
| | switch(config-role)# no description | Resets the description for the Tech. Docs. group. |

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed to perform configuration commands, and role2 users are only allowed to perform debug commands, then if Joe belongs to both role1 and role2, he can perform configuration as well as debug commands.



Note

If you belong to multiple roles, you can execute a superset 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 Rules and Features for Each Role

The **rule** command specifies rules for a specific group using the order in which they were issued. The order is important because the order in which they were issued is the order in which they are applied.

The **feature** option specifies the features accessible for this rule (for example, FSPF, zone, VSAN, fcping, interface). If no features are specified, the rule applies to all **show**, **debug**, **clear**, configuration commands, and all **exec** commands.

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

In this case, **exec** commands refer to all commands in the EXEC mode that do not fall in the **show**, **debug**, and **clear**, categories. Up to 16 rules can be configured for each role.

To configure a new 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 sangroup switch(config-role)# | Places you in role mode for the sangroup role submode. |
| 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 <i>sangroup</i> role to perform all configuration commands except fspf 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 <i>sangroup</i> to perform the fcping command. |

In Step 3, rule 1 is applied first, thus permitting all **config** commands to sangroup users. 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 since the second rule globally overrode the first rule.

Displaying Role-Based CLI Information

Use the **show role** command to display rules configured on the switch including those rules that have not yet been committed to persistent storage. The rules are displayed by rule number and are based on each role. All roles are displayed even if role name is not specified. See [Example 14-3](#).

Example 14-3 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: sangroup
Description: SAN management group
-----
Rule   Type   Command-type   Feature
-----
1.    permit   config         *
2.     deny   config         fspf
3.    permit   debug         zone
4.    permit   exec          fcping
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring CLI User Profiles

Every Cisco MDS 9000 Family switch user has related NMS 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 CLI commands explained in this section enable you to create users and modify the profile of an existing user. These commands are restricted to privileged users as determined by your administrator.

Creating or Updating Users

The switches use the same command (**username**) to create a user and to update an existing user. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format. By default, the user account does not expire unless you explicitly configure it to expire.



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, nsd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys



Note

User passwords are not displayed in the switch configuration file.

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 abcd expire 2003-05-31 switch(config)# | Creates or updates the user account (usam) along with a password (abcd) that is set to expire on 2003-05-31. The password is limited to 64 characters. |
| | switch(config)# username msam password 0 abcd role network-operator switch(config)# | Creates or updates the user account (msam) along with a password (abcd) specified in clear text (indicated by 0). The password is limited to 64 characters. |
| | switch(config)# username user1 password 5 !@*asdsfsdfjh!@df switch(config)# | Specifies an encrypted (specified by 5) password (!@*asdsfsdfjh!@df) for the user account (user1). |
| Step 3 | switch(config)# username usam role network-admin switch(config)# | Adds the specified user (usam) to the network-admin role. |
| | switch(config)# no username usam role vsan-admin switch(config)# | Deletes the specified user (usam) from the vsan-admin role. |



Note

If the **update-snmpv3** option is used, specify the clear text and old SNMP password (see the “[Forcing Identical SNMP and CLI Passwords](#)” section on page 14-24).

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying User Profile Information

Use the **show user-account** command to display configured information about user accounts. See Examples 14-4 to 14-6.

Example 14-4 Displays All Users

```
switch# show users
switch# show users
admin    pts/7          Jan 12 20:56 (10.77.202.149)
admin    pts/9          Jan 12 23:29 (modena.cisco.com)
admin    pts/10         Jan 13 03:05 (dhcp-171-71-58-120.cisco.com)
admin    pts/11         Jan 13 01:53 (dhcp-171-71-49-49.cisco.com)
```

Example 14-5 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 14-6 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring CLI Accounting Parameters

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 purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

Setting the Accounting Log Size

The **aaa accounting logsize** command sets the size limit of the accounting log file in persistent storage. The default is 15,000 bytes.

To set the log file size, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# aaa accounting logsize 29000 | Sets the size of the log file on the local disk. The default is 15,000 bytes. |

Enabling RADIUS Accounting

To enable RADIUS accounting in a switch, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# aaa accounting method radius | Configures the RADIUS accounting method on the switch. By default, only local accounting is enabled. |

You can clear the RADIUS accounting configuration by issuing the **no aaa accounting method radius** command.



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* flag in the AAA client configuration. This flag should be turned on 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Accounting Configuration

The **show accounting** command displays configured accounting information. See Examples 14-7 to 14-9.

Example 14-7 Displays Configured Accounting Parameters.

```
switch# show accounting config
RADIUS accounting not enabled
local accounting enabled
```

Example 14-8 Displays Configured Log Size.

```
switch# show accounting logsize
maximum local accounting log size:29000
```

Example 14-9 Displays the Entire Log File.

```
switch# show accounting log
Tue Jan 15 06:03:24 1980:update:::Created interface vsan1

Tue Jan 15 06:15:08 1980:start:/dev/pts/0_316764908:admin
Tue Jan 15 07:26:10 1980:stop:/dev/pts/0_316764908:admin:vsh exited normally
Tue Jan 15 07:46:40 1980:update:/dev/ttyS0_316753046:admin:Alias test is created
on VSAN 1

Tue Jan 15 07:46:40 1980:update:/dev/ttyS0_316753046:admin:Alias test is removed
on VSAN 1

Tue Jan 15 08:01:57 1980:update:/dev/ttyS0_316753046:admin:in-order delivery gua
rantee settings changed in-order guarantee:yes
Tue Jan 15 08:02:31 1980:update:/dev/ttyS0_316753046:admin:Enabled IP routing

Tue Jan 15 08:09:51 1980:update:/dev/ttyS0_316753046:admin:Alias test is created
on VSAN 1

Tue Jan 15 08:09:52 1980:update:/dev/ttyS0_316753046:admin:Alias test is removed
on VSAN 1

Tue Jan 15 08:11:30 1980:update:/dev/ttyS0_316753046:admin:Zone test is created
on VSAN 1
.
.
.
Sat Jan 19 22:16:06 1980:update:/dev/pts/0_317167691:admin:Zone cisco is removed
on VSAN 1

Sat Jan 19 22:56:49 1980:stop:/dev/pts/0_317167691:admin:vsh exited normally
Sun Jan 20 17:07:50 1980:start:snmp_317236070_10.77.202.149:public
Sun Jan 20 17:07:50 1980:stop:snmp_317236070_10.77.202.149:public:
Mon Jan 21 03:38:03 1980:start:/dev/pts/0_317273883:admin
Mon Jan 21 04:08:25 1980:stop:/dev/pts/0_317273883:admin:vsh exited normally
Mon Jan 21 06:43:49 1980:start:/dev/pts/0_317285029:admin
Mon Jan 21 06:44:38 1980:stop:/dev/pts/0_317285029:admin:vsh exited normally
Mon Jan 21 07:24:16 1980:start:/dev/pts/0_317287456:admin
```

Send documentation comments to mdsfeedback-doc@cisco.com

Recovering Administrator Password

An administrator can recover a password from a local console connection.

The password recovery procedure must be performed on the supervisor module that becomes the active supervisor module after the recovery procedure is completed. To ensure the other supervisor module does not become the active module, you have two options:

- Physically remove the other supervisor module from the chassis, or
- Change the other supervisor module's console prompt to the `loader>` or `switch(boot)#` prompt (see the [“Recovery from the loader> Prompt”](#) section on page 5-30) until you complete this procedure.



Note

Password recovery is not possible from a Telnet or SSH session.

To recover a administrator's password, follow these steps:

Step 1 Reboot the switch.

```
switch# reload
The supervisor is going down for reboot NOW!
```

Step 2 Press the **Ctrl-]** key sequence (when the switch begins its SAN-OS software boot sequence) to enter the `switch(boot)#` prompt (see [“Recovery Interruption”](#) section on page 5-26).

```
Ctrl-]
switch(boot)#
```

Step 3 Change to configuration mode.

```
switch(boot)# config terminal
```

Step 4 Enter the **admin-password** command to reset the administrator password.

```
switch(boot-config)# admin-password password
```

Step 5 Exit to the EXEC mode.

```
switch(boot-config)# exit
switchboot#
```

Step 6 Enter the **load** command to load the SAN-OS software.

```
switch(boot)# load bootflash:system.img
```

Step 7 Save the software configuration.

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

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring RADIUS Authentication

You can configure RADIUS query parameters. These commands are restricted to privileged users as determined by your administrator.

Setting the RADIUS Server Address

You can add up to five (5) RADIUS servers using the **radius-server host** command. You can configure a RADIUS server to be a primary server so it is always contacted first. If you have not configured a primary server, the RADIUS servers are tried in the order they were configured.

To specify the RADIUS server address and the 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 primary switch(config)# | Adds 10.10.0.0 users to the RADIUS server list as the primary server. This server is always tried first. |
| Step 3 | switch(config)# radius-server host 10.10.0.0 key HostKey switch(config)# | Specifies a 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 4 | switch(config)# radius-server host 10.10.0.0 auth-port 2003 switch(config)# | 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 5 | switch(config)# radius-server host 10.10.0.0 acct-port 2004 switch(config)# | 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 6 | switch(config)# radius-server host 10.10.0.0 accounting switch(config)# | Specifies this server to be used only for accounting purposes. Note If neither the authentication option nor the accounting options are specified, the server is used for both accounting and authentication purposes. |
| Step 7 | switch(config)# radius-server host radius1 primary switch(config)# | Specifies the server to be the primary server. |
| | switch(config)# radius-server host radius2 key 0 abcd switch(config)# | Specifies a clear text key for the specified server. The key is restricted to 65 characters. |
| | switch(config)# radius-server host radius3 key 7 1234 switch(config)# | Specifies a reversible encrypted key for the specified server. The key is restricted to 65 characters. |

Send documentation comments to mdsfeedback-doc@cisco.com

Setting the RADIUS 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 65 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.

To set the RADIUS preshared key, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server key AnyWord switch(config)# | 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 switch(config)# | 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 public switch(config)# | Configures a preshared key (public) specified in encrypted text (indicated by 7) to authenticate communication between the RADIUS client and server. |

Setting the RADIUS Server Time-Out Interval

To specify the time 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 switch(config)# | Specifies the time (in seconds) between retransmissions to the RADIUS server. The default time-out is one (1) second. The time range in seconds is 1 to 60. |

You can revert the retransmission time to its default by issuing the **no radius-server timeout** command.

Send documentation comments to mdsfeedback-doc@cisco.com

Setting Iterations of the RADIUS Server

By default, a switch retries a RADIUS server connection only once. This number can be configured. The maximum is five retries per server. You can revert the retry number to its default by issuing the **no radius-server retransmit** command.

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 switch(config)# | Configures the number of times (3) the switch tries to connect to a RADIUS server(s) before reverting to local authentication. |

The worst case cumulative response or timeout latency from RADIUS servers for authentication should not be more than 50 sec. For example in the following configuration:

```
radius-server timeout 5
radius-server retransmit 3
radius-server host A authentication
radius-server host B authentication
```

The worst case cumulative response or timeout latency will be:

```
(5+1)*3    +    (5+1)*3 = 36
^^^^^^^^  ^^^^^^^  ^^^^^
server A    server B    total
```



Note

You need to add one (1) to the retransmit count to calculate the total. The total number of tries equals the number of retransmits + 1.

Defining 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 sep value *
```

Where `protocol` is a Cisco attribute for a particular type of authorization, and `sep` is = for mandatory attributes, and * 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, like authorization information, along with authentication results. This authorization information is specified through VSAs.

Send documentation comments to mdsfeedback-doc@cisco.com

VSA Format

The following VSA protocol options are supported:

- 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:

- `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. This is an example using the roles attribute:

```
Cisco-AVPair = "shell: roles = "network-admin vsan-admin" "
```

- `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 value.

Authorization Process

The RADIUS based authorization process is as follows:

-
- | | |
|---------------|--|
| Step 1 | The switch sends an Access-Request packet to the RADIUS server. |
| Step 2 | <p>The RADIUS server responds with an Accept or Reject message.</p> <ul style="list-style-type: none">• If Access-Reject is received, that means authentication has failed and no authorization information is sent.• If Access-Accept is received, that means authentication is successful and VSA is also sent along with the Access-Accept packet.• If no VSA data is sent, local authorization is used.• If your user name has no corresponding local account, a new account is created. This new account is locked and cannot be used for local login. It is deleted after 24 hours. |
| Step 3 | You are made a member of all groups indicated in the role list attribute in the VSA. You are removed from those roles if your role is not listed in the VSA group list. |
-

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Displaying RADIUS Server Details

Use the **show radius-server** command to display all configured RADIUS server parameters (see [Example 14-10](#)).



Note

Only administrators can view the RADIUS preshared key.

Example 14-10 Displays Configured RADIUS Information

```
switch# show radius-server
Global RADIUS shared secret:Myxggc
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:23MHcUnD
  10.10.0.0:
    available for authentication on port:1812
    available for accounting on port:1813
    RADIUS shared secret:hostkey----> for administrators only
```

Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a host key pair. To generate a host key, use the **ssh key** command (see the [“Generating an SSH Host Key Pair”](#) section on page 14-19).

Enabling SSH Service

By default, the SSH service is disabled. To enable SSH service, issue the **ssh server enable** command. 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. |


Send documentation comments to mdsfeedback-doc@cisco.com

Generating an SSH Host Key Pair

Be sure to have an SSH host key pair with the appropriate version before enabling the SSH service. The SSH service accepts three types of key pairs for use by SSH versions 1 and 2. Generate the SSH host key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

- 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.

To generate the SSH host 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 switch(config)# | Generates the RSA1 host key pair. |
| | switch(config)# ssh key dsa 1024 generating dsa key..... generated dsa key switch(config)# | Generates the DSA host key pair. |
| | switch(config)# ssh key rsa 1024 generating rsa key..... generated rsa key switch(config)# | Generates the RSA host key pair. |
| | switch(config)# no ssh key rsa 1024 cleared RSA keys switch(config)# | Clears the RSA host key pair configuration. |
| | |  Caution If you delete all of the SSH keys, you cannot start a new session. |

Using the force Option

If the SSH key pair option is already generated for the required version, use the **force** option 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 switch(config)# | Tries to set the host key pair. If a required host key pair is already configured, use the force option to overwrite that host key pair. |
| | | Deletes the old DSA key and sets the host key pair using the new bit specification. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 14-11](#).

Example 14-11 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 host key pair details for the specified key or for all keys, if no key is specified. See [Example 14-12](#).

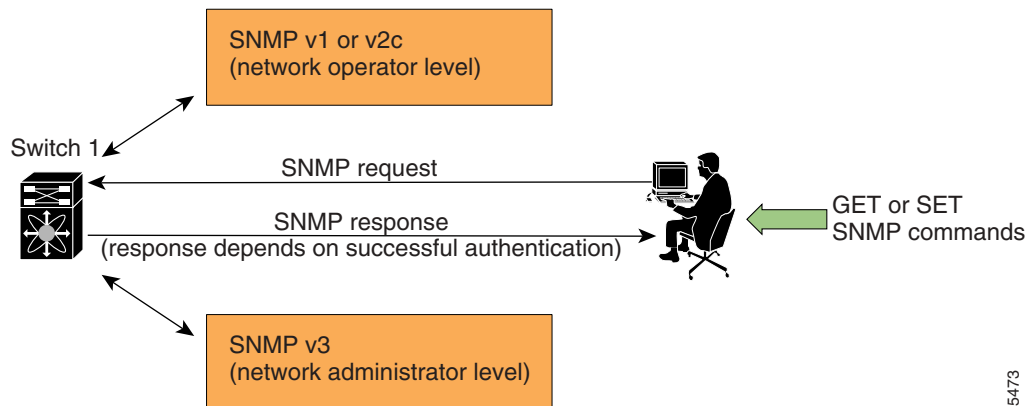
Example 14-12 Displays Host Key Pair Details

```
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+OTn8eGdnnDVxw5eJs5OcOEXOyjaW
cMMYsEgxc9ada1NElp8Wy7GPMWGOQYj9CU0AAAAMCWhNN18zFNOIPo7cU3t7d0iEbAAAAQBdQ8UAO
i/Cti84qFb3kTqXLS9mEhdQUo0lHcH5bw5PKfj2Y/dLR437zCBKXetPj4p7mhQ6Fq5os8RZtJEyOsNsA
AABAA0oxZbPyWer5NHATXiyXdPI7j9i8fgyn9FNipMkOF2Mn75Mi/lqQ4NIq0gQNvQ0x27uCeQlRts/Q
wI4q68/eaw==
fingerprint:
512 f7:cc:90:3d:f5:8a:a9:ca:48:76:9f:f8:6e:71:d4:ae
```

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 14-4](#)).

Figure 14-4 SNMP Security



85473

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

SNMP users are different from CLI users. SNMP users also have role-based authentication for roles and authorization purposes.

SNMP Version 1 and Version 2c

SNMPv1 and 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

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.

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.

Switches in the Cisco MDS 9000 Family are preconfigured with two default groups: network-admin and network-operator. The network-admin role has the access rights to define custom roles with custom privileges. You can begin communicating with the agent once the your user name is created, your roles are set up by your administrator, and you are added to the roles.

**Note**

Users and roles configured through the CLI are different from users and groups configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Restricting Switch Access

You can restrict SNMPv1 or SNMPv2c access to a Cisco MDS 9000 Family switch. This restriction can only be performed using SNMP directly. You can not use the Cisco MDS 9000 Family Fabric Manager or CLI to perform this restriction.



Note

This restriction is possible only for SNMPv1 or SNMPv2c access.

To restrict access to an MDS switch using SNMP, follow these steps.

Step 1

Create an entry for each hosts that required access in the `snmpTargetAddressTable` with the same `snmpTargetAddressTagList` (for example, `list100`).



Note

All other fields in the `snmpTargetAddressTable` are not shown in this example to focus attention on the fields required for this step.

| <code>snmpTargetAddrTAddress</code> | <code>snmpTargetAddrTagList</code> |
|-------------------------------------|------------------------------------|
| 11.71.49.67 | list100 |
| 11.71.11.54 | list100 |

Step 2

For each community name in the `snmpCommunityTable` for which access to the switch is required, you must set the `snmpCommunityTransportTag` value to match the `snmpTargetAddressTagList` value.

The following example displays the `snmpCommunityTable` if access is restricted to the *public* community:

| <code>snmpCommunityName</code> | <code>snmpCommunityTransportTag</code> |
|--------------------------------|--|
| mycommunity | " " |
| yourcommunity | " " |
| public | list100 |



Note

If the `snmpCommunityTransportTag` has a zero-length string, the access list is not applied.

Creating SNMP Groups

You can create SNMP groups through SNMP by configuring a group entry in the `vacmAccessTable` on the switch.

Refer to RFC2575.

Send documentation comments to mdsfeedback-doc@cisco.com

Creating and Modifying Users

You can create users or modify existing users using SNMP or the CLI.

- **SNMP**—Create a user as a clone of an existing user in the `vsmUserTable` on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC2574.



Note You must explicitly configure password(s) for SNMP users. The SNMP user passwords are not generated as the part of the configuration file as they are not portable across devices. The password is limited to a minimum of 8 characters and a maximum of 64 characters.



Tip An SNMP user must be created on each switch to which the user requires access. If the user is managing 10 switches, each of the 10 switches must have the SNMP user defined.

- **CLI**—You can create a user or modify an existing user using the **`snmp-server user`** command.

By default only two groups are available in a Cisco MDS 9000 Family switch—`network-operator` and `network-admin`. To assign a user to a new SNMP group, you must first create the SNMP group.



Note SNMP groups can only be created in SNMP (see the [“Creating SNMP Groups” section on page 14-22](#)).

To create or modify SNMP users using the CLI, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# snmp-server user joe network-admin auth sha abcd1234</code> | Creates or modifies the settings for a user (joe) in the network-admin group using the HMAC-SHA-96 authentication password (abcd1234). |
| | <code>switch(config)# snmp-server user sam network-admin auth md5 abcdefgh</code> <code>switch(config)#</code> | Creates or modifies the settings for a user (sam) in the network-admin group using the HMAC-MD5-96 authentication password (abcdefgh). |
| | <code>switch(config)# snmp-server user Bill network-admin auth sha abcd1234 priv abcdefgh</code> | Creates or modifies the settings for a user (network-admin) in the network-admin group using the HMAC-SHA-96 authentication level and privacy encryption parameters. |
| | <code>switch(config)# no snmp-server user usernameA</code> | Deletes the user (usernameA) and all associated parameters. |
| | <code>switch(config)# snmp-server user user1 network-admin auth md5 0xab0211gh priv 0x45abf342 localizedkey</code> | Specifies the password to be in localized key format (see RFC2574). The localized key is provided in the hex format (for example, 0xacbdef). |



Note Avoid using the **`localizedkey`** option when configuring an SNMP user from 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Forcing Identical SNMP and CLI Passwords

You can force the SNMPv3 password and the CLI password to be the same. You must know the SNMPv3 password to change the password using the CLI. Use CLI password to synchronize the SNMP password. The password is limited to a minimum of 8 characters and a maximum of 64 characters.



Caution

To change the SNMP password, a clear text CLI password is required.

To modify the secret key for an SNMPv3 user, refer to RFC2574.

To update the SNMPv3 password from the CLI, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# username joe password wxyz6789 update-snmpv3 abcd1234 | Updates the SNMPv3 password for the specified user (joe). The local CLI password and the SNMP password are updated. If user Joe does not exist, the command fails. |

Assigning Users to Groups

Once the user and the group are created, the administrator should configure an entry in the vacmSecurityToGroupTable to add the configured user to a configured group.

To assign users to groups through SNMP, refer to RFC2575.

To assign users to groups through the CLI, refer to the procedure specified in the [“Creating and Modifying Users”](#) section on page 14-23.

Adding or Deleting Communities

You can configure read-only or read-write access for SNMP users by using the **snmp-server community** CLI command. Use the **no** form of the command to delete the configured community. Refer to RFC2576.

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). |

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying SNMP Security Information

Use the **show snmp** commands to display configured SNMP information (see [Example 14-13](#) and [14-15](#)).

Example 14-13 Displays SNMP User Details

```
switch# show snmp user
```

| User | Group | Auth | Priv |
|--------|------------------|------|------|
| steve | network-admin | md5 | des |
| sadmin | network-admin | md5 | des |
| stever | network-operator | md5 | des |

Example 14-14 Displays SNMP Community Information

```
switch# show snmp community
```

| Community | Access |
|------------|--------|
| private | rw |
| public | ro |
| v93RACqPNH | ro |

Example 14-15 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 |
| ... | | | | | |

Displaying SNMP Counter Information

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 (see the *Cisco MDS 9000 Family Fabric Manager User Guide*). See [Example 14-16](#).

Example 14-16 Displays SNMP

```
switch# show snmp
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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

0 Too big errors
1 No such name errors
0 Bad values errors
0 General errors
Community
-----
public
User
-----
admin

Access
-----
rw
Group
-----
network-admin

Auth Priv
-----
md5 no

```

Default Security Settings

Table 14-2 lists the default settings for all security features in any switch.

Table 14-2 Default Security Settings

| Parameters | Default |
|--|---|
| Roles in each switch (for CLI users) | Two default roles—network-operator and network-admin. |
| Group in each switch (for SNMP users) | Two default roles—network-operator and network-admin. |
| AAA authentication login | Local authentication is enabled. If the Telnet or SSH options are not specified, the command applies to both. |
| Telnet server | Enabled. |
| Accounting log file size on local disk | 15,000 bytes. |
| User's account expiration | Does not expire unless you explicitly configure it to expire. |
| RADIUS server timeout interval | The default time-out is five (5) seconds. |
| RADIUS preshared key | No key is configured. |
| RADIUS server connection attempts | A switch tries to connect to a RADIUS server up to 3 times. |
| RADIUS Authentication messages | 1812 messages sent by destination UDP port. |
| RADIUS Accounting messages | 1813 messages sent by destination UDP port. |
| User name | admin. |
| User password | admin. |



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.
 - FSPF 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:

- [FSPF Features, page 15-2](#)
- [FSPF Examples, page 15-2](#)
- [Configuring FSPF Globally, page 15-3](#)
- [Configuring FSPF for a Specific Interface, page 15-5](#)
- [Configuring Fibre Channel Routes, page 15-7](#)
- [Clearing FSPF Counters, page 15-8](#)
- [Broadcast Routing, page 15-9](#)
- [In-Order Delivery, page 15-9](#)
- [Configuring Flow Statistics, page 15-12](#)
- [Displaying Routing and Forwarding Information, page 15-14](#)
- [Default Settings, page 15-18](#)

Send documentation comments to mdsfeedback-doc@cisco.com

FSPF Features

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 15-1 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 15-1 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).


Send documentation comments to mdsfeedback-doc@cisco.com

Redundant Links

To further improve on the topology in Figure 15-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 15-2 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 15-2 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.

Configuring FSPF Globally

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you don't 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.

To configure a FSPF feature for the entire VSAN, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fspf config vsan 1 switch-config-(fspf-config)# | Enters FSPF global configuration mode for the specified VSAN. |
| Step 3 | switch-config-(fspf-config)# spf static switch-config-(fspf-config)# | Forces static SPF computation for the dynamic (default) incremental VSAN. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|--|--|
| Step 4 | switch-config-(fspf-config)# spf hold-time 10 switch-config-(fspf-config)# | Configures the hold time between two route computations in milliseconds 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 switch-config-(fspf-config)# | Configures the autonomous region for this VSAN and specifies the region ID (7). |

Deleting the Entire FSPF Configuration

To delete FSPF configuration for the entire VSAN, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf config vsan 3 switch(config)# | Deletes the FSPF configuration for VSAN 3. |

Disabling FSPF Routing Protocols

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

To enable or disable FSPF routing protocols, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf enable vsan 5 switch(config)# | Disables FSPF routing protocol in VSAN 5. |
| | switch(config)# fspf enable vsan 7 switch(config)# | Enables FSPF routing protocol in VSAN 7. |

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 15-1](#) displays the default settings for switch responses.

Table 15-1 LSR Default Settings

| LSR Option | Default | Description |
|---|------------|--|
| Acknowledgement interval (RxmtInterval) | 5 seconds | The time a switch waits for an acknowledgement 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring FSPF for a Specific Interface

Several FSPF commands are available on a per interface basis. The following configuration procedures apply to an interface in a specific VSAN and are described in this section.

- [Computing Route Cost, page 15-5](#)
- [Specifying Hello Time Intervals, page 15-5](#)
- [Specifying Dead Intervals, page 15-6](#)
- [Disabling FSPF for Specific Interfaces, page 15-6](#)
- [Retransmitting Intervals, page 15-7](#)

Computing Route 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 1Gbps is 1000 and 2Gbps is 500.

To configure FSPF link cost, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# confi g t | 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 switch(config-if)# | Configures the cost for the selected interface in VSAN 90. |

Specifying 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.

To configure the FSPF Hello time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# confi g t | 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Specifying Dead 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.

To configure the FSPF dead time interval, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | 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. |

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.

To disable FSPF for a specific interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

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.


To configure the FSPF retransmit time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | 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. |

Configuring Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. To configure 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 15-3](#)).

Figure 15-3 Fibre Channel Routes



Note

Other than in VSANs, run time checks are not performed on configured and suspended static routes.

Send documentation comments to mdsfeedback-doc@cisco.com

To configure an FC route, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| 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: ff0000 matches only the domain, ffff00 matches the domain and the area, ffffff matches the domain, area, and port. |

1. FIB = Forwarding Information Base

Clearing FSPF Counters

To clear the FSPF statistics counters for one interface or for the entire VSAN, follow this step:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# clear fspf counters vsan 1 switch# | Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared. |
| | switch# clear fspf counters vsan 200 interface fc1/1 switch# | Clears the FSPF statistics counters for the specified interface in VSAN 200. |

Send documentation comments to mdsfeedback-doc@cisco.com

Broadcast Routing

Broadcast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric (for broadcast traffic).

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 the distribution tree information. The protocols create a loop-free broadcast distribution tree.



Caution

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

In-Order Delivery

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

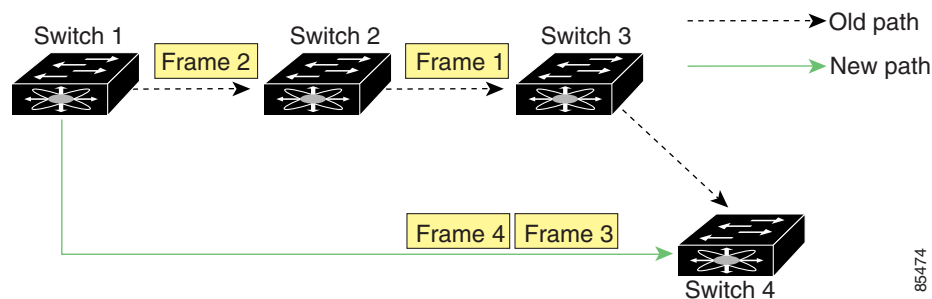
Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

In case of a single switch, 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.

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 (see [Figure 15-4](#)).

Figure 15-4 Route Change Delivery



In [Figure 15-4](#), the new path from Switch 1 to Switch 4 is faster. Hence, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

Send documentation comments to mdsfeedback-doc@cisco.com


If the in-order guarantee feature is enabled, the frames within the network are treated as specified below:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames which can not be delivered in-order within the network latency drop period are dropped inside the network.
- The number of dropped frames are reduced by slowing down the traffic at the frame source.

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 (see [Figure 15-5](#)).

Figure 15-5 Link Congestion Delivery



In [Figure 15-5](#), the port of the old path (red dot) is congested. Hence Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order guarantee feature is enabled, the frames crossing a PortChannel are treated as specified below:

- Frames using the old path are delivered before new frames are accepted.
- Frames which cannot be delivered in-order, through the old path, within the switch latency drop period are dropped.
- The new frames are delivered through the new path after the switch latency drop period has elapsed.

Enabling In-Order Delivery

By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.



Tip

We recommend that you only enable this feature in a switch when devices are present in the switch that cannot handle any out-of-order frames. 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 due to an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out of order.

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

```
switch# show in-order-guarantee
inorder delivery is not guaranteed
```

Send documentation comments to mdsfeedback-doc@cisco.com

To enable in-order delivery, 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. |

Configuring the Drop Latency Time

Use this command if you need to change the default latency time for either a network or a 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)# fcdroplateny network 5000 | Configures network drop latency time to be 5000 milliseconds for the network. The valid range is 0 to 60000 milliseconds. The default is 2000 milliseconds. Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network |
| | switch(config)# no fcdroplateny network 4500 | Removes the current fcdroplateny network configuration (4500) and reverts the switch to the factory defaults. |
| Step 3 | switch(config)# fcdroplateny switch 4000 | Configures switch drop latency time to be 4000 milliseconds for the switch. The valid range is 0 to 60000 milliseconds. The default is 500 milliseconds. Note The switch drop latency parameter should have the same value in all the switches in the network |
| | switch(config)# no fcdroplateny switch 4500 | Removes the current fcdroplateny switch configuration (4500) and reverts the switch to the factory defaults. |

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplateny** command (see [Example 15-1](#)).

Example 15-1 *Displays Administrative Distance*

```
switch# show fcdroplateny
switch latency value:4000 milliseconds
network latency value:5000 milliseconds
```

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring Flow Statistics

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

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics. 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.

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. |

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 ffffff vsan 1 switch(config)# | Enables the flow counter. Note The source ID and the destination ID are specified in the FC ID hex format (for example, 0x123aff). The mask can be one of ff0000 or ffffff. |
| Step 3 | switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2 switch(config)# | Disables the flow counter. |

Send documentation comments to mdsfeedback-doc@cisco.com

Clearing FIB¹ Statistics

To clear the aggregated flow counter, use the **clear fcflow stats** command (see Examples 15-2 and 15-3).

Example 15-2 Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Example 15-3 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 15-4 to 15-6).

Example 15-4 Displays Aggregated fcflow 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 15-5 Displays fcflow 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 15-6 Displays fcflow Index Usage for the Specified Module

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

1. FIB = Forwarding Information Base

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying Routing and Forwarding Information

You can view specific information about existing Fibre Channel and FSPF configurations at any time from the EXEC mode. The following **show** commands provide further details on existing Fibre Channel paths and routes (see Examples 15-7 to 15-15).



Note

When the number of routes are displayed in the command output, both visible and hidden routes are include in the total number of routes. While the hidden routes are added to the count, they will not be visible.

Example 15-7 Displays Administrative Distance

```
switch# show fcroute distance
```

| UUID | Route Distance | Name |
|------|-------------------|------------|
| ---- | ----- | ---- |
| 10 | 20 | RIB |
| 22 | 40 | FCDOMAIN |
| 39 | 80 | RIB-CONFIG |
| 12 | 100 | FSPF |
| 17 | 120 | FLOGI |
| 21 | 140 | TLPM |
| 14 | 180 | MCAST |
| 64 | 200 | RIB-TEST |

Example 15-8 Displays Multicast Routing Information

```
switch# show fcroute multicast
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| ---- | ----- | ----- |
| 1 | 0xffffffff | 0 |
| 2 | 0xffffffff | 1 |
| 3 | 0xffffffff | 1 |
| 4 | 0xffffffff | 0 |
| 5 | 0xffffffff | 0 |
| 6 | 0xffffffff | 0 |
| 7 | 0xffffffff | 0 |
| 8 | 0xffffffff | 0 |
| 9 | 0xffffffff | 0 |
| 10 | 0xffffffff | 0 |

Example 15-9 Displays FCID Information for a Specified VSAN

```
switch# show fcroute multicast vsan 3
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| ---- | ----- | ----- |
| 3 | 0xffffffff | 1 |

Example 15-10 Displays FCID and interface Information for a Specified VSAN

```
switch# show fcroute multicast 0xffffffff vsan 2
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| ---- | ----- | ----- |
| 2 | 0xffffffff | 1 |
| | fc1/1 | |

Send documentation comments to mdsfeedback-doc@cisco.com

Example 15-11 Displays Unicast Routing Information

```
switch# show fcroute unicast
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  1    0x010101 0xffffffff 0x00 0x00 D P A 1    10
static  2    0x111211 0xffffffff 0x00 0x00 R P A 1    10
fspf    3    0x610000 0xff0000 0x00 0x00 D P A 4    500
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 15-12 Displays Unicast Routing Information for a Specified VSAN

```
switch# show fcroute unicast vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 15-13 Displays Unicast Routing Information for a Specified FCID

```
switch# show fcroute unicast 0x040101 0xffffffff vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
      fc1/2 Domain 0xa6(166)
```

Example 15-14 Displays Route Database Information

```
switch# show fcroute summary
FC Route Database Created Thu Feb 13 07:21:52 2003
VSAN      Ucast      Mcast      Label      Last Modified Time
-----
1          5          1          0          Thu Feb 13 10:21:06 2003
2          4          1          0          Thu Feb 13 10:21:07 2003
3          4          1          0          Thu Feb 13 10:21:08 2003
4          4          1          0          Thu Feb 13 10:21:09 2003
5          4          1          0          Thu Feb 13 10:21:10 2003
6          4          1          0          Thu Feb 13 10:21:11 2003
7          4          1          0          Thu Feb 13 10:21:12 2003
8          4          1          0          Thu Feb 13 10:21:13 2003
9          4          1          0          Thu Feb 13 10:21:14 2003
10         4          1          0          Thu Feb 13 10:21:15 2003
11         4          1          0          Thu Feb 13 10:21:16 2003
12         4          1          0          Thu Feb 13 10:21:17 2003
13         4          1          0          Thu Feb 13 10:21:18 2003
14         4          1          0          Thu Feb 13 10:21:18 2003
15         4          1          0          Thu Feb 13 10:21:19 2003
-----
Total      61          15          0
```

Send documentation comments to mdsfeedback-doc@cisco.com

Example 15-15 Displays Route Database Information for a Specified VSAN

```
switch# show fcroute summary vsan 5
FC Route Database Created Thu Feb 13 07:21:52 2003
```

| VSAN | Ucast | Mcast | Label | Last Modified Time |
|-------|-------|-------|-------|--------------------------|
| 5 | 4 | 1 | 0 | Thu Feb 13 10:21:10 2003 |
| Total | 4 | 1 | 0 | |

Displaying Global FSPF Information

The **show fspf** command (see [Example 15-16](#)) displays global FSPF information for a specific VSAN:

- the domain number of the switch
- the autonomous region for the switch
- Min_LS_arrival: the minimum time that must elapse before the switch accepts LSR updates
- Min_LS_interval: the minimum time that must elapse before the switch can transmit an LSR
- LS_refresh_time: the interval lapse between refresh LSR transmissions
- Max_age: the maximum time aa LSR can stay before being deleted

Example 15-16 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

The **show fspf database** command displays a summary of the FSPF database for a specified VSAN (see [Example 15-17](#)). 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

Send documentation comments to mdsfeedback-doc@cisco.com

- LSR age
- LS 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 15-17 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 |

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying FSPF Interfaces

The **show fspf** command displays the following information for each selected interface (see [Example 15-18](#)).

- 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 15-18 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 15-2](#) lists the default settings for FSPF features.

Table 15-2 Default FSPF Settings

| Parameters | Default |
|---|---|
| FSPF | Enabled on all E ports and TE ports. |
| SPF computation | Dynamic. |
| SPF hold time | 0. |
| Backbone region | 0. |
| Acknowledgement interval (RxmtInterval) | 5 seconds. |
| Refresh time (LSRefreshTime) | 30 minutes. |
| Maximum age (MaxAge) | 60 minutes. |
| Hello interval | 20 seconds. |
| Dead interval | 80 seconds. |
| Distribution tree information | Derived from the principal switch (root node). |
| Routing table | FSPF stores up to 16 equal cost paths to a given destination. |

Send documentation comments to mdsfeedback-doc@cisco.com

Table 15-2 Default FSPF Settings (continued)

| Parameters | Default |
|---------------------------|---|
| 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. |

Send documentation comments to mdsfeedback-doc@cisco.com



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 (NMS):

- 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.

This chapter includes the following sections:

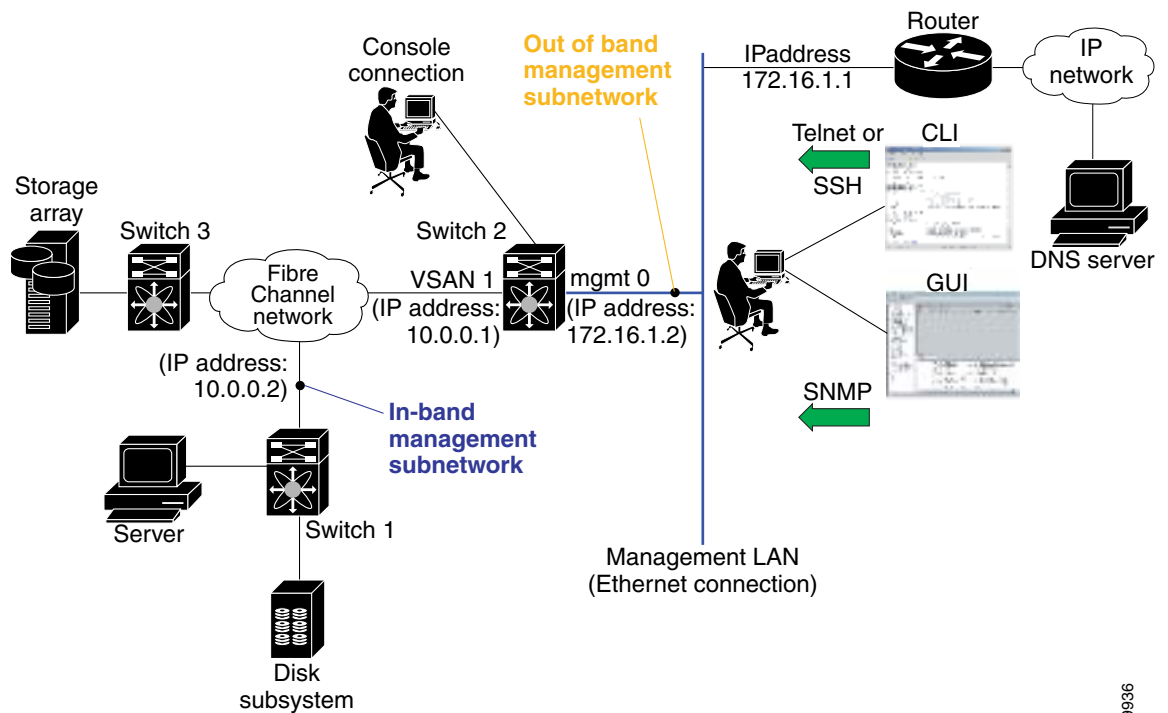
- [Traffic Management Services, page 16-2](#)
- [Configuring the Ethernet Management Port, page 16-2](#)
- [Configuring the Default Gateway, page 16-3](#)
- [Configuring the Default Network, page 16-4](#)
- [Configuring IPFC, page 16-5](#)
- [Configuring IP Static Routes, page 16-6](#)
- [Displaying IP Interface Information, page 16-7](#)
- [Configuring Overlay VSANs, page 16-8](#)
- [Configuring Multiple VSANs, page 16-10](#)
- [Configuring VRRP, page 16-12](#)
- [Configuring DNS Server, page 16-19](#)
- [Default Settings, page 16-20](#)

Send documentation comments to mdsfeedback-doc@cisco.com

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running IP protocol over a 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 16-1).

Figure 16-1 Management Access to Switches



79936

Configuring the Ethernet Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP network management sessions. You can also configure the supervisor module's Ethernet interface and VSAN interfaces as management ports. This section focuses on the Ethernet management port (mgmt0). You can remotely configure the switch through the management port. To configure a connection remotely, you must configure the IP parameters (IP address and subnet mask) from the CLI so that the switch is reachable.



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.

Send documentation comments to mdsfeedback-doc@cisco.com

To configure the mgmt0 Ethernet interface, 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 1.1.1.1 255.255.255.0 | Enters the IP address (1.1.1.1) and IP subnet mask (255.255.255.0) for the management interface. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |

Configuring the Default Gateway

Use the **IP default-gateway** command to configure the IP address for a switch's default gateway. This IP address should be configured along with the IP 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 IP addresses will be used regardless of the default-gateway being enabled or disabled. If these IP address are configured and 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 3-2).

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 IP address of the gateway switch (see [Figure 16-2](#)).

Figure 16-2 Overlay VSAN Functionality



In [Figure 16-2](#), switch A has the IP address 1.1.2.11.1, switch B has the IP address 1.1.2.11.2, switch C has the IP address 1.1.2.11.3, and switch D has the IP address 1.1.2.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.10 to connect to the gateway switch.

Send documentation comments to mdsfeedback-doc@cisco.com

Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IP 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 [“Configuring VSAN Interfaces”](#) section on page 9-15).

To configure default gateways, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip default-gateway 1.12.11.1 switch(config)# | Configures the IP address for the default gateway (1.12.11.1). |

Use the **show ip route** command to verify that the IP address for the default gateway is configured.

Configuring the Default Network

Unlike the **ip default-gateway** command, use the **ip default-network** command when IP routing is enabled on the switch. If you assign the IP default network address, the switch considers routes to that network as the last resort. If the IP default network address is not available, the switch uses the IP default gateway address. For every network configured with the IP 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 IP addresses will be used regardless of the default gateway being enabled or disabled. If these IP address are configured and 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 3-2).

To configure default networks, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip default-network 190.10.1.0 switch(config)# 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 switch(config)# | Configures the IP address for the default network (190.10.1.0). Defines a static route to network 10.0.0.0 as the static default route. |

Use the **show ip route** command to verify if the IP address for the default gateway is configured.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN using the **ip address** command. If you wish to override the default IP Address, use the **ip address** command.

Configuring an IP Address in a VSAN

To configure a VSAN interface and an IP address for that interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures the interface for the specified VSAN (1). |
| Step 3 | switch(config-if)# ip address 10.0.0.12 255.255.255.0 switch(config-if)# | Configures the IP address and netmask for the selected interface. |

Disabling IP Routing

By default, the IP routing feature is disabled in all switches. To enable the IP routing feature, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip routing switch(config)# | Enables IP routing (disabled by default). |
| Step 3 | switch(config)# no ip routing switch(config)# | Disables IP routing and reverts to the factory settings. |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring IP Static Routes

Static routing is a mechanism to configure IP routes on the switch. You can configure more than one static route.

To configure a static route, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# IP route <network IP address> <netmask> <next hop IP 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 IP address, subnet mask, next hop, and distance, and VSAN or management interface. |

If your configuration does not need an external router, you can use static routing.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IP 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.

Viewing and Clearing ARPs

Address Resolution Protocol (ARP) entries can be viewed (**show arp**), deleted (**no arp**), or cleared (**clear arp-cache**) in Cisco MDS 9000 Family switches. 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
switch(config)#
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
switch#
```

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying IP Interface Information

Use the following **show** commands to view configured IP interface information (see Examples 16-1 to 16-4).

Example 16-1 Displays the VSAN Interface

```
switch# show interface vsan1
vsan1 is up, line protocol is up
  WWPN is 10:00:00:05:30:00:59:1f, FCID is 0x9c0100
  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
```



Note

You can see the output for this command only if you have previously configured a virtual network interface (see the “Configuring an IP Address in a VSAN” section on page 16-5).

Example 16-2 Displays the Connected and Static Route Details

```
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 16-3 Displays Configured Routes

```
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 |

Example 16-4 Displays the IP Routing Status

```
switch# show ip routing
ip routing is disabled
```

Send documentation comments to mdsfeedback-doc@cisco.com

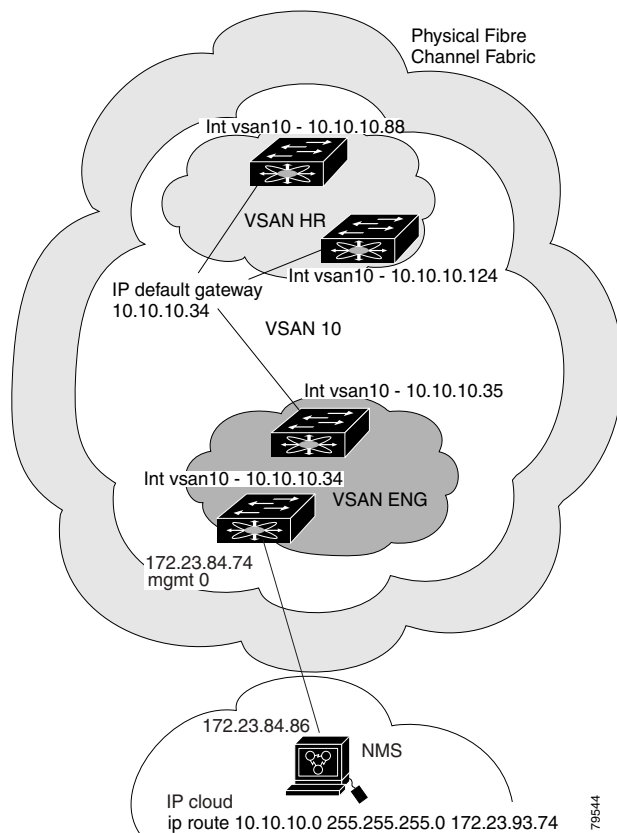
Configuring Overlay VSANs

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.

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 switch 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 default gateway (route) and the IP address on switches that point to the NMS (see [Figure 16-3](#)).

Figure 16-3 Overlay VSAN Configuration Example



Send documentation comments to mdsfeedback-doc@cisco.com

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 16-3](#)), 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 switch--config-vsan-db# | 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 vsan10 switch(config-if)# | Creates a VSAN interface (VSAN 10). |
| Step 6 | switch(config-if)# ip address 10.10.10.x netmask 255.255.255.0 switch(config-if)# | Assigns an IP address and netmask on all switches in the fabric. |
| Step 7 | switch(config-if)# no shut | 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 16-3](#), follow this step:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# 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. |



Note

To configure the management interface displayed in [Figure 16-3](#), set the default gateway to an IP address on the Ethernet network.

Send documentation comments to mdsfeedback-doc@cisco.com

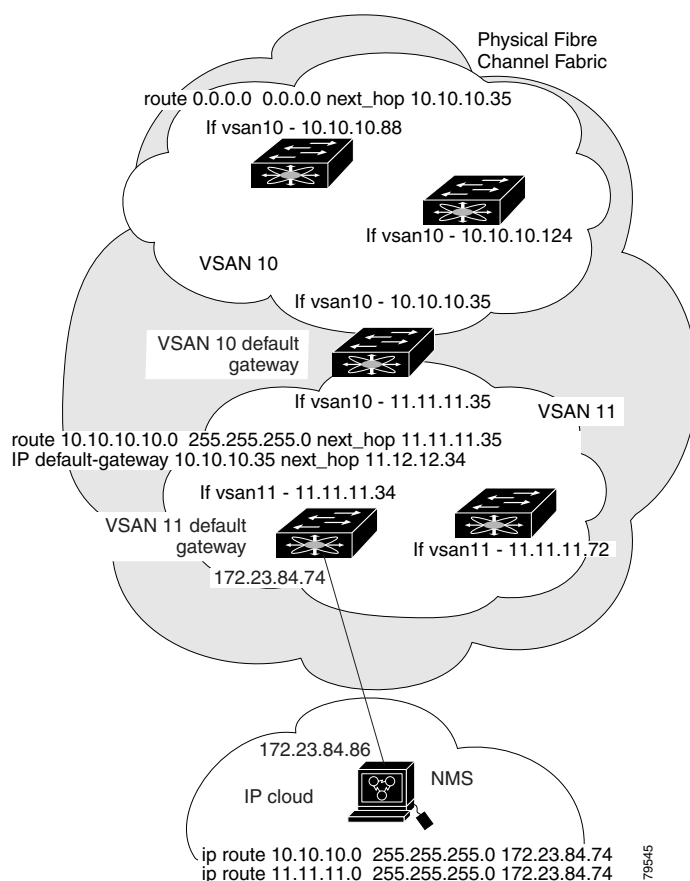
Configuring Multiple VSANs

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 an overlay VSAN, 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 route on the Fibre Channel switches and the IP cloud (see [Figure 16-4](#)).

Figure 16-4 Multiple VSANs Configuration Example



Send documentation comments to mdsfeedback-doc@cisco.com

To configure an overlay VSAN (using the example in Figure 16-4), 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_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 database 10 mode. |
| 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. |
| Step 6 | switch-config-vsan-db# exit switch(config)# | Exits the VSAN database 11 mode. |
| Step 7 | switch(config)# interface vsan10 switch(config-if)# | Enters the VSAN 10 interface configuration mode for VSAN 10. |
| Step 8 | switch(config-if)# ip address 10.10.10.x netmask 255.255.255.0 switch(config-if)# | Assigns an IP address and netmask on all switches in VSAN 10. |
| Step 9 | switch(config-if)# no shut | 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 vsan11 switch(config-if)# | Enters the VSAN 11 interface configuration mode. |
| Step 12 | switch(config-if)# ip address 11.11.11.x netmask 255.255.255.0 switch(config-if)# | Assigns an IP address and netmask on all of the switches in VSAN 11. |
| Step 13 | switch(config-if)# no shut | 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 IP 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring VRRP

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.

VRRP Features

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.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with a different virtual router IP mapping.
- 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.

Send documentation comments to mdsfeedback-doc@cisco.com

VRRP Functionality

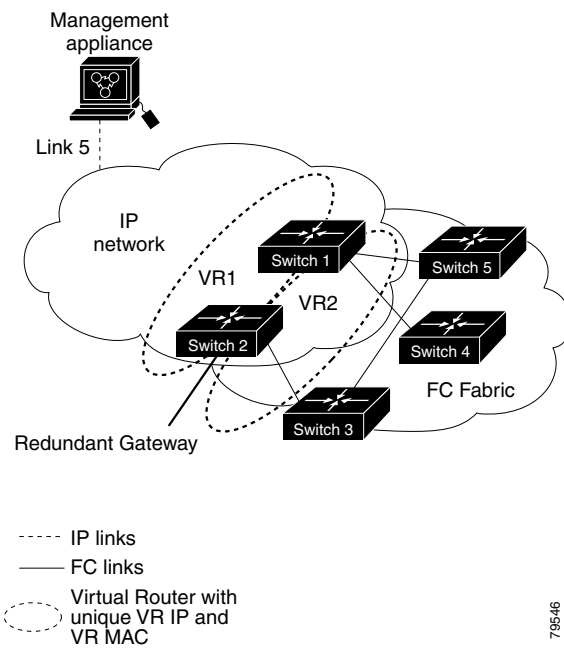
In [Figure 16-5](#), switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches don't have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 16-5 VRRP Functionality



In [Figure 16-6](#), the fabric example has two virtual router groups (VR1 and VR2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR1 and the FC interface is in VR2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 16-6 Redundant Gateway



[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Creating or Removing a Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

To create or remove a VR, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# config t switch(config)# | 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 VR ID 250. |
| | switch(config-if-vrrp)# no vrrp 250 switch(config-if) | Removes a VR ID 250. |

Enabling a Virtual Router

By default, a virtual router is always disabled (**shutdown**). VRRP can be configured only if this state is disabled. Be sure to configure at least one IP address before attempting to enable a VR.

To enable or disable a virtual router, follow these steps:

| | Command | Purpose |
|--------|---|------------------------------|
| Step 1 | switch(config-if-vrrp)# no shutdown switch(config-if-vrrp)# | Enables VRRP configuration. |
| | switch(config-if-vrrp)# shutdown switch(config-if-vrrp)# | Disables VRRP configuration. |

Adding an IP Address for a Virtual Router

One primary IP address and multiple secondary addresses can be configured for a switch. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address.

To configure an IP address 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)# interface ipaddress 10.0.0.12 address switch(config-if)# | Configures an IP address. The IP address must be configured before the VRRP is added. |
| Step 4 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates VR ID 250. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|--|--|
| Step 5 | switch(config-if-vrrp)# address 10.0.0.10 switch(config-if-vrrp)# | Configures the IP address (10.0.0.10) for the selected VR. |
| | switch(config-if-vrrp)# no address 10.0.0.10 switch(config-if-vrrp)# | Removes the IP address (10.0.0.10) for the selected VR. |

Setting 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 a switch with the primary IP address.

To set the priority 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)# priority 2 switch(config-if-vrrp)# | Configures the priority for the selected VRRP. |
| | | Note Priority 255 cannot be preempted. |

Setting the Time Interval for the Advertisement Packet

The valid time range for an advertisement packet is between 1 and 255 seconds with the default being 1 (one) second. If the switch has the primary IP address, this time must be specified.

To set the priority 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)# advertisement-interval 15 switch(config-if-vrrp)# | Sets the interval time in seconds between sending advertisement frames. |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Preempting the Master Virtual Router

By default, the preempt option is enabled. An owner with priority 255 cannot be preempted. If two priorities match, the owner with the highest priority preempts the master virtual router.

To enable or disable preempting, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# confi g 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)# preempt switch(config-if-vrrp)# | 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 switch(config-if-vrrp)# | Disables the preempt option and allows the master to keep its priority level. |

Configuring Authentication for the Virtual Router

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

To set an authentication option for a virtual router, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# confi g 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|--------|--|---|
| Step 4 | switch(config-if-vrrp)# authentication text password switch(config-if-vrrp)# | Assigns the simple text authentication option and specifies the password for this option. |
| | switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003 switch(config-if-vrrp)# | Assigns 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 switch(config-if-vrrp)# | Assigns the no authentication option, which is the default. |

Setting the Priority Based on Interface State

The tracking feature is disabled by default. When you specify the tracking option, the priority of the virtual router is changed based on the state of another interface in the switch. When the tracked interface is down, the priority of the virtual router is changed to a lower priority value. When the tracked interface is up, the priority of the virtual router is restored to its original value. You can track one of two interfaces on a switch in the Cisco MDS 9000 Family: a specified VSAN interface or a management interface.

To track the interface priority 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)# track interface mgmt 0 priority 2 switch(config-if-vrrp)# | Specifies the priority of the virtual router to be modified based on the state of the management interface. |
| | switch(config-if-vrrp)# no track switch(config-if-vrrp)# | Disables the tracking feature. |

Displaying VRRP Information

Use the **show vrrp vr** command to display configured VRRP information (see Examples 16-5 to 16-8).

Example 16-5 Displays 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Example 16-6 Displays 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 16-7 Displays 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
```

Example 16-8 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** command to clear all the software counters for the specified virtual router (see [Example 16-9](#)).

Example 16-9 Clears VRRP Information

```
switch# clear vrrp 7 interface vsan2
switch#
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring DNS Server

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

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-address translation and reverts to the factory default. |
| Step 3 | switch(config)# no ip domain-name cisco.com | Disables the domain name and reverts to the factory default. |
| | switch(config)# ip domain-name cisco.com | Enables (default) 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. |
| 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 list of default domain names to complete unqualified host names, use the ip domain-list global configuration command. You can define up to 10 domain names in this list. To delete a name from a list, use the no form of this command. |
| | switch(config)# no ip domain-list | Deletes the defined list 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 did configure 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 15.2.0.0 | Specifies the first address (15.1.0.1) as the primary server and the second address (15.2.0.0) as the secondary sever. 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. |
| 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. | | |

The DNS server may be dropped after two attempts due to the following reasons:

- if the IP address or the switch name is wrongly configured
- if the DNS server is not reachable due to 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see [Example 16-10](#)).

Example 16-10 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 16-1](#) lists the default settings for IP features.

Table 16-1 Default IPFC Settings

| Parameters | Default |
|--|---------------------------------------|
| VSAN IP interface configuration | No IP address is assigned by default. |
| IP routing | Disabled. |
| Domain lookup | Disabled. |
| Domain name | Enabled. |
| Domain list | No domains are configured. |
| Name server | No servers are configured. |
| Virtual router | Disabled (shutdown). |
| Virtual router priority for switches with secondary IP address | 100. |
| Virtual router priority for switches with primary IP address | 255. |
| Time interval between advertisement frames | 1 second. |
| Preempting master VR | Enabled. |
| VRRP security authentication | No authentication. |
| Interface state tracking | Disabled. |



Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services modules 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 allows IP hosts to access Fibre Channel storage using iSCSI protocol.

This chapter includes the following sections:

- [IP Storage Services Module, page 17-2](#)
- [Configuring Gigabit Ethernet Interfaces, page 17-4](#)
- [Configuring FCIP, page 17-17](#)
- [Configuring iSCSI, page 17-38](#)
- [Default IP Storage Settings, page 17-76](#)

Send documentation comments to mdsfeedback-doc@cisco.com


IP Storage Services Module

The IPS services module (IPS module) allows you to use FCIP and iSCSI features. It integrates seamlessly into the Cisco MDS 9000 Family, and supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports. Each port can run FCIP and iSCSI 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 17-1](#) depicts the FCIP scenarios in which the IPS module is used.

Figure 17-1 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 MDS 9000 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 17-2](#) depicts the iSCSI scenarios in which the IPS module is used.

Figure 17-2 iSCSI Scenarios



Send documentation comments to mdsfeedback-doc@cisco.com

Verifying the Module Status

After inserting the module, verify the status of the module using the **show module** command:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    16     1/2 Gbps FC Module        DS-X9016             ok
4    8      IP Storage Module        DS-X9308-SMIP       ok <-----IPS module
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)
---  -
2    1.1(1)      0.3         20:41:00:05:30:00:86:5e to 20:50:00:05:30:00:86:5e
4    1.1(1)      0.2         20:c1:00:05:30:00:86:5e to 20:c8:00:05:30:00:86:5e
5    1.1(1)      0.602      --
6    1.1(1)      0.602      --

Mod  MAC-Address(es)                Serial-Num
---  -
2    00-05-30-00-9f-62 to 00-05-30-00-9f-66  JAB064505YV
4    00-05-30-00-a1-ae to 00-05-30-00-a1-ba  JAB0649059h
5    00-05-30-00-9f-f6 to 00-05-30-00-9f-fa  JAB06350B1M
6    00-05-30-00-9f-f2 to 00-05-30-00-9f-f6  JAB06350B1F

* this terminal session
```

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring Gigabit Ethernet Interfaces

This section includes the following topics:

- [About Gigabit Ethernet Interfaces, page 17-4](#)
- [Basic Gigabit Ethernet Configuration, page 17-4](#)
- [About VLANs for Gigabit Ethernet, page 17-5](#)
- [VLAN Configuration, page 17-6](#)
- [Interface Subnet Requirements, page 17-6](#)
- [Managing IP Routing, page 17-7](#)
- [Verifying Gigabit Ethernet Connectivity, page 17-7](#)
- [Managing ARP Caches, page 17-8](#)
- [Displaying Statistics, page 17-8](#)
- [Gigabit Ethernet High Availability, page 17-13](#)
- [Configuring CDP, page 17-16](#)
- [IPS Core Dumps, page 17-16](#)

About Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On the IPS 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.

A new port mode, called **IPS**, is defined for Gigabit Ethernet ports on the IPS 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.




Tip

Gigabit Ethernet ports on the IPS 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

Figure 17-3 depicts a basic Gigabit Ethernet configuration.

Figure 17-3 Gigabit Ethernet Configuration



Send documentation comments to mdsfeedback-doc@cisco.com

To configure the Gigabit Ethernet interface for the scenario in [Figure 17-3](#), 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 (slot2, port 2). |
| Step 3 | switch(config-if)# ip address 10.1.1.100 255.255.255.0 | Enters the IP 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. |

You can configure the switch to receive and transfer large (or jumbo) frames on a port. The default IP MTU frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased to 9000 bytes. In this example, the size was set to 3000 bytes. Independent of the MTU size, the IPS module does not pack multiple IP frames (converted to FCIP or to iSCSI).



Note

The minimum MTU size for a port running iSCSI is 620 bytes.

To configure MTU frame size, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config-if)# switchport mtu 3000 switch(config-if)# | Changes the IP maximum transmission unit (MTU) to 3000. The default is 1500. |

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.

The Gigabit Ethernet port can be configured as a trunking port and uses the IEEE 802.1Q standard for VLAN encapsulation.



Note

If the IPS module is connected to a Cisco Ethernet switch, verify the following requirements:

- the Ethernet switch port is configured as a trunking port, and
- the encapsulation is set to 802.1Q and not ISL, which is the default.

Send documentation comments to mdsfeedback-doc@cisco.com

VLAN Configuration

To configure a VLAN subinterface (the 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 (slot2, 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 10.1.1.100 255.255.255.0 | Enters the IP address (10.1.1.100) and IP 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 interface (major), subinterfaces (VLAN tags) and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 17-1](#)).

Table 17-1 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 tag 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 tags cannot be configured in the same subnet. |
| Gigabit Ethernet 1/1 | Gigabit Ethernet 1/1.100 | No | A VLAN tag 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 17-1](#) also applies to Ethernet PortChannels.

Send documentation comments to mdsfeedback-doc@cisco.com

Managing IP Routing

To configure static IP routing 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 IP address of the router connected to the Gigabit Ethernet interface. |

Displaying the IP Route Table

The **show ips ip route interface ethernet** command takes the ethernet interface as a parameter and returns the route table for the interface. See [Example 17-1](#).

Example 17-1 Displays the 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.

Verifying Gigabit Ethernet Connectivity

The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the “[Using the ping Command](#)” section on page 2-18).

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch using the **ping** command. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly. See [Example 17-2](#).

Example 17-2 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
```



Note

If the connection fails, verify the following, and repeat the **ping** command:

- 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, and
- the Gigabit Ethernet interface is in the **up** state (use the **show interface gigabitethernet** command).

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Managing ARP Caches

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 17-3](#).

Example 17-3 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
...
```

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache. See Examples [17-4](#) and [17-5](#).

Example 17-4 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 17-5 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

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 Gigabit Ethernet** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 17-6](#).

Example 17-6 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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 17-7 Displays the Gigabit Ethernet's 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 17-8](#).

Example 17-8 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 Ethernet statistics for that interface. See [Example 17-9](#).

Example 17-9 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

```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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/IP Statistics



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

Use the **show ips stats ip interface gigabitethernet** to display and verify IP statistics. This command takes the main Gigabit Ethernet interface as a parameter and returns IP statistics for that interface. See [Example 17-10](#).

Example 17-10 Displays IP 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

```

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 [Examples 17-11](#) and [17-12](#).

Example 17-11 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 |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Example 17-12 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 17-13](#).

Example 17-13 Displays ICMP Statistics

```
switch# show ips stats icmp interface gigabitethernet 4/1
ICMP Statistics for port GigabitEthernet4/1
  5 ICMP messages received
  0 ICMP messages dropped due to errors
ICMP input histogram
  5 echo request
ICMP output histogram
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
5 echo reply
```


Send documentation comments to mdsfeedback-doc@cisco.com

Gigabit Ethernet High Availability

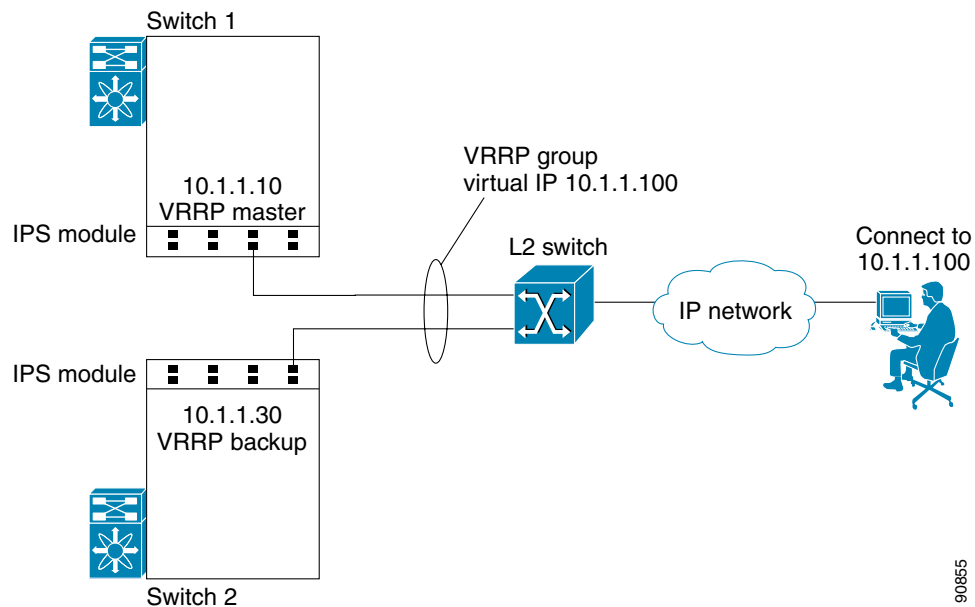
Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

Configuring VRRP

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services (see the “Configuring VRRP” section on page 16-12).

VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 17-4).

Figure 17-4 VRRP Scenario



In Figure 17-4, 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
- Interfaces across IPS modules in one switch
- Interfaces across IPS modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels
- Subinterfaces

Send documentation comments to mdsfeedback-doc@cisco.com

To configure VRRP for Gigabit Ethernet interfaces, 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 (slot2, port 2). |
| Step 3 | switch(config-if)# ip address 10.1.1.10 255.255.255.0 | Enters the IP address (10.1.1.10) and IP 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 a VR ID 100. |
| Step 6 | switch(config-if-vrrp)# address 10.1.1.100 | Configures the virtual IP address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IP address must be in the same subnet as the IP address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IP 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. |



Note

The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Configuring Ethernet PortChannels

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.

The data traffic from one TCP connection always travels on the same physical links. An Ethernet switch connecting to the MDS Gigabit Ethernet port can implement load balancing based on its IP address, its source-destination MAC address, or its IP and port. If Ethernet-based load balancing cannot be implemented for iSCSI scenarios based on the IP and port, multiple iSCSI initiators are required to take advantage of the Ethernet PortChannel feature.



Note

The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3aa protocol.


Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 17-5](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

**Note**

PortChannel members must be one of these combinations: ports 1-2, ports 3-4, ports 5-6, or ports 7-8.

Figure 17-5 Ethernet PortChannel Scenario



In [Figure 17-5](#), Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel.

**Note**

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that FCIP link.

PortChannel configuration specified in [Chapter 11, “Configuring PortChannels”](#) also apply to Ethernet PortChannel configurations.

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.

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 IP address (10.1.1.1) and IP 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). |
| 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

| | Command | Purpose |
|---------|---|--|
| Step 7 | <code>switch(config-if)# no shutdown</code> | Enables the selected interface. |
| Step 8 | <code>switch(config)# interface gigabitethernet 9/4</code> <code>switch(config-if)#</code> | Configures the specified Gigabit Ethernet interface (slot 9, port 4). |
| Step 9 | <code>switch(config-if)# channel-group 10</code> <code>gigabitethernet 9/4 added to port-channel 10</code> 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 | <code>switch(config-if)# no shutdown</code> | Enables the selected interface. |



Note

Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:

- if the interface already has an IP address assigned, or
- if subinterfaces are configured on that interface.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. See the [“Configuring CDP” section on page 3-33](#).

IPS 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 sometimes useful to obtain a full 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. Core dumps take up significant space and hence the level of what gets stored can be configured using one of the two options:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files).
- Full core dumps—Each full core dump consists of 75 parts (75 files). This dump cannot be saved on the supervisor module due to its large space requirement. If you choose this option, then you must configure an external TFTP server using the **system cores tftp:** command (see [“Configuring Core and Log Files” section on page 26-6](#)).

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring FCIP

This section includes the following topics:

- [About FCIP, page 17-17](#)
- [Basic FCIP Configuration, page 17-20](#)
- [Advanced FCIP Profile Configuration, page 17-22](#)
- [Advanced FCIP Interface Configuration, page 17-26](#)
- [E Port Configurations, page 17-32](#)
- [Displaying FCIP Information, page 17-33](#)
- [FCIP High Availability, page 17-35](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 17-37](#)

About FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed 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 17-6](#).

Figure 17-6 Fibre Channel SANs Connected by FCIP

FCIP uses Transmission Control Protocol (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).

To configure the IPS module for FCIP, you should have a basic understanding of the following concepts:

- [FCIP and VE Ports, page 17-18](#)
- [FCIP Link, page 17-18](#)
- [FCIP Profiles, page 17-19](#)
- [FCIP Interface, page 17-19](#)

Send documentation comments to mdsfeedback-doc@cisco.com

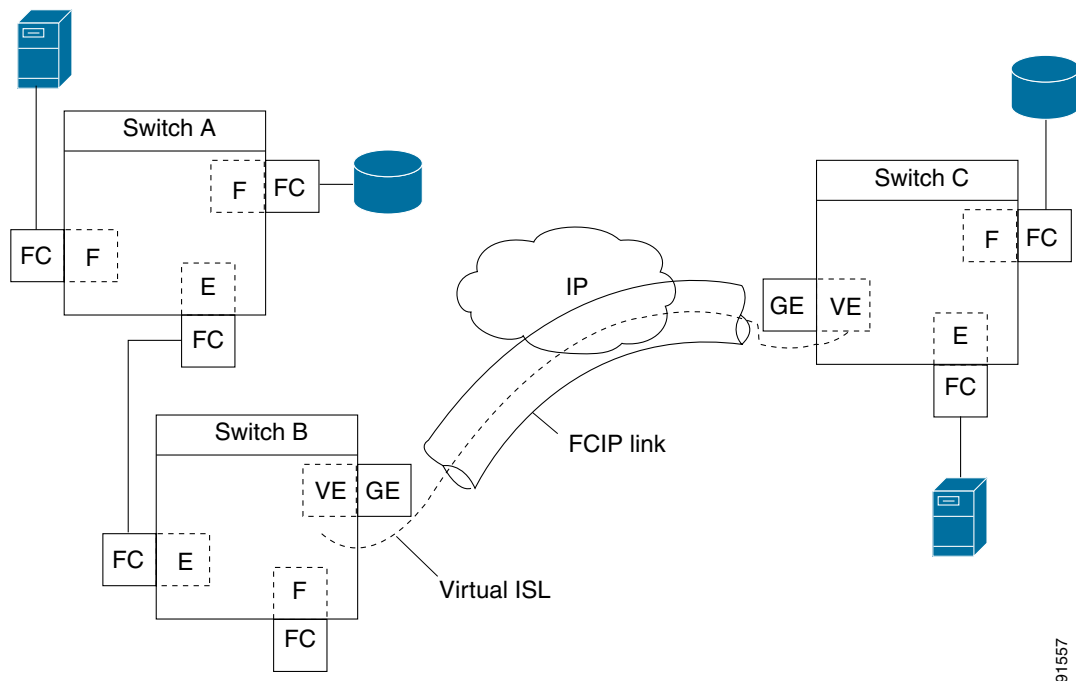
FCIP and VE Ports

Figure 17-7 provides the internal model of FCIP with respect to Fibre Channel inter switch links (ISLs) and Cisco's enhanced ISLs (EISLs). See the “E Port” section on page 9-3.

FCIP defines virtual E (VE) ports, which 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 a 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 17-7).

Figure 17-7 FCIP Links and Virtual ISLs



91557

FCIP Link

FCIP links consist of one or more TCP connections between two FCIP link end points. 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 second connection is used only for Fibre Channel control frames, i.e. switch-to-switch protocol frames (all Class F) frames. This arrangement is used to provide low latency for all control frames.

To enable FCIP on the IPS module, a FCIP profile and FCIP interface (interface FCIP) must be configured.

Send documentation comments to mdsfeedback-doc@cisco.com

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based on 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-switch 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 local IP address and TCP parameters. The profile defines 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 terminates (see [Figure 17-8](#)).

Figure 17-8 FCIP Profile and FCIP Links

FCIP Interface

The FCIP interface is the local end point of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port terminates FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

Send documentation comments to mdsfeedback-doc@cisco.com

Basic FCIP Configuration


To configure a FCIP link, perform this procedure on both switches.

- Step 1** Configure the Gigabit Ethernet interface.
- Step 2** Create a FCIP profile, assign the Gigabit Ethernet interface's IP address to the profile. See the [“Creating FCIP Profiles”](#) section on page 17-20.
- Step 3** Create a FCIP interface, assign the profile to the interface. See the [“Creating FCIP Links”](#) section on page 17-21.
- Step 4** Configure the peer IP address for the FCIP interface. See the [“Creating FCIP Links”](#) section on page 17-21.
- Step 5** Enable the interface. See the [“Creating FCIP Links”](#) section on page 17-21.

Creating FCIP Profiles

To create a FCIP profile, you must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile (see [Figure 17-9](#)).

Figure 17-9 Assigning Profiles to Each Gigabit Ethernet Interface



To create a FCIP profile in switch 1, 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 IP address of the Gigabit Ethernet interface (3/1). |

To assign FCIP profile in switch 2, 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 IP address of the Gigabit Ethernet interface. |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Creating FCIP Links

When two FCIP link end points are created, a FCIP link is established between the two IPS modules. To create a FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) a FCIP link to that peer switch (see Figure 17-10).

Figure 17-10 Assigning Profiles to Each Gigabit Ethernet Interface



To create a FCIP link end point 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 a FCIP interface (51). |
| Step 3 | switch1(config-if)# use-profile 10 | Assigns the profile (10) to the FCIP interface. |
| Step 4 | switch1(config-if)# peer-info ipaddr 10.1.1.1 | Assigns the peer IP 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 a FCIP link end point in switch 2, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch2# config terminal switch2(config)# | Enters configuration mode. |
| Step 2 | switch2(config)# interface fcip 52 switch2(config-if)# | Creates a 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 address 10.100.1.25 | Assigns the peer IP address information (10.100.1.25 for switch 1) to the FCIP interface |
| Step 5 | switch1(config-if)# no shutdown | Enables the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

- [Configuring TCP Listener Ports, page 17-22](#)
- [Configuring TCP Parameters, page 17-22](#)

FCIP configuration options can be accessed from the `switch(config-profile)#` submode prompt.

To enter the `switch(config-profile)#` prompt, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# fcip profile 20</code> <code>switch(config-profile)#</code> | Creates the profile (if it does not already exist). The valid range is from 1 to 255. |

Configuring TCP Listener Ports

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 | <code>switch(config-profile)# port 5000</code> | Associates the profile with the local port number (5000). |
| | <code>switch(config-profile)# no port</code> | Reverts to the default 3225 port. |

Configuring TCP Parameters

This section provides details on the TCP parameters that can be configured to control TCP behavior in a switch. The following TCP parameters can be configured.

- [Minimum Retransmit Timeout, page 17-23](#)
- [Keepalive Timeout, page 17-23](#)
- [Maximum Retransmissions, page 17-23](#)
- [Path MTU, page 17-24](#)
- [SACK, page 17-24](#)
- [Window Management, page 17-24](#)
- [Buffer Size, page 17-25](#)
- [Quality of Service, page 17-25](#)
- [Monitoring Window Congestion, page 17-26](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Minimum Retransmit Timeout

The **tcp minimum-retransmit-time** option controls the minimum amount of time TCP waits before retransmitting. By default, this value is 300 milliseconds.

To configure the minimum retransmit time, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-profile)# tcp min-retransmit-time 500</code> | Specifies the minimum TCP retransmit time for the TCP connection in milliseconds (500). The default is 300 milliseconds and the range is from 250 to 5000 milliseconds. |
| | <code>switch(config-profile)# no tcp min-retransmit-time 500</code> | Reverts the minimum TCP retransmit time to the factory default of 300 milliseconds. |

Keepalive Timeout

The **tcp keepalive-timeout** option enables you to configure the interval between which the TCP connection verifies if the FCIP link is functioning. This ensures that a FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified transmission time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to detect FCIP link failures.

The first interval during which the connection is idle is 60 seconds (default). When the connection is idle for 60 seconds, 8 keepalive probes are sent at 1-second intervals. If no response is received for these 8 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 from the default of 60 seconds. This interval is identified using the **keepalive-timeout** option. The valid range is from 1 to 7200 seconds.

To configure the keep alive timeout, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-profile)# tcp keepalive-timeout 120</code> | Specifies the keepalive timeout interval for the TCP connection in seconds (120). The default is 60 seconds. The range is from 1 to 7200 seconds. |
| | <code>switch(config-profile)# no tcp keepalive-timeout 120</code> | Reverts the keepalive-timeout to 60 seconds. |

Maximum Retransmissions

The **tcp max-retransmissions** option specifies 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 | <code>switch(config-profile)# tcp max-retransmissions 6</code> | Specifies the maximum number of retransmissions (6). The default is 4 and the range is from 1 to 8 retransmissions. |
| | <code>switch(config-profile)# no tcp max-retransmissions 6</code> | Reverts to the default of 4 retransmissions. |

Send documentation comments to mdsfeedback-doc@cisco.com

Path MTU

Path MTU (PMTU) is the minimum MTU on the IP network between the two end points 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 default timeout of 3600 seconds. If TCP reduces the size of the max 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 | <code>switch(config-profile)# no tcp pmtu-enable</code> | Disables PMTU discovery. |
| | <code>switch(config-profile)# tcp pmtu-enable</code> | Enables (default) PMTU discovery with the default value of 3600 seconds. |
| | <code>switch(config-profile)# tcp pmtu-enable reset-timeout 90</code> | Specifies the PMTU reset timeout to 90 seconds. The default is 3600 seconds and the range is from 60 to 3600 seconds. |
| | <code>switch(config-profile)# no tcp pmtu-enable reset-timeout 600</code> | Leaves the PMTU in an enabled state but changes the timeout to the default of 3600 seconds. |

SACK

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 time. 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 | <code>switch(config-profile)# no tcp sack-enable</code> | Disables SACK. |
| | <code>switch(config-profile)# tcp sack-enable</code> | Enables SACK (default). |

Window Management

The optimal TCP window size is computed using three options.

- The **maximum-bandwidth** option configures the maximum available end-to-end bandwidth in the path (900 Mbps in the configuration example).
- The **minimum-available-bandwidth** option configures the minimum slow start threshold.
- The **round-trip-time** option is the estimated round trip time across the IP network to reach the FCIP peer end point (10 milliseconds in the configuration example). If the round-trip-time value is under-estimated, the TCP window size will be too small to reach the maximum available bandwidth. If the round-trip-time is overestimated, the TCP window size will be too big. If the maximum available bandwidth is correct, this will cause increase in latency and potential packet drop in the network but will not affect the speed.

The **maximum-bandwidth** option and the **round-trip-time** option together determine the window size.

Send documentation comments to mdsfeedback-doc@cisco.com

The **minimum-available-bandwidth** option and the **round-trip-time** option together determine the threshold below which TCP aggressively increases its size. After it reaches the threshold the software uses standard TCP rules to reach the maximum available bandwidth. The defaults are max-bandwidth = 1G, min-available-bandwidth = 2 Mbps, and round-trip-time is 10ms

To configure window management, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>switch(config-profile)# tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code> | Configures the maximum available bandwidth at 900 Mbps, the minimum slow start threshold as 300 Mbps, and the round trip time as 10 milliseconds. |
| | <code>switch(config-profile)# no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code> | Reverts to the factory defaults. The defaults are max-bandwidth = 1G, min-available-bandwidth = 2 Mbps and round-trip-time is 10ms. |
| | <code>switch(config-profile)# tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200</code> | Configures the maximum available bandwidth at 2000 Kbps, the minimum slow start threshold as 2000 Kbps, and the round trip time as 200 microseconds. |

Buffer Size

The **send-buffer-size** option defines 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 buffer size is 0 KB.

To set the buffer size, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>switch(config-profile)# tcp send-buffer-size 5000</code> | Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 8192 KB. |
| | <code>switch(config-profile)# no tcp send-buffer-size 5000</code> | Reverts the switch to its factory default (0 KB). |

Quality of Service

The **qos control** option 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 connection.

To set the control values, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config-profile)# tcp qos control 3 data 5</code> | Configures the control TCP connection and data connection to mark all packets on that DSCP value. |
| | <code>switch(config-profile)# no tcp qos control 3 data 5</code> | Reverts the switch to its factory default (no packets). |

Send documentation comments to mdsfeedback-doc@cisco.com

Monitoring Window Congestion

By configuring the congestion window monitoring (CWM) option, you can influence the rate at which TCP ramps up the transmitted bandwidth after an idle period as listed below:

- If the traffic is transmitted in burst sizes that are smaller than the configured CWM value, you can send the whole traffic burst immediately, provided no drops occurred.
- If the traffic burst is larger than the configured CWM value, some traffic will not be sent immediately.
- If the end-to-end path between the two Cisco MDS 9000 Family switches is 1G, you can set the maximum burst size.
- If the router connecting to the IPS port does not have sufficient buffering, you can use the smallest available value to decrease the burst size.

By default the **tcp cwm** option is enabled and the default burst size is 10KB.



Tip

We recommend that this feature remain enabled to realize optimal performance.

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 defaults burst size at 10 KB. |
| | 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 cp cwm burstsize 25 | Leaves the CWM feature in an enabled state but changes the burst size to the default of 10 KB. |

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface. To do so, you must first create the interface and enter the `config-if` submode.

- [Configuring Peers, page 17-27](#)
- [Configuring Active Connection, page 17-28](#)
- [Configuring the Number of TCP Connections, page 17-29](#)
- [Enabling Time Stamps, page 17-29](#)
- [B Port Interoperability Mode, page 17-30](#)

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 a FCIP interface (100). |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring Peers

To establish a FCIP link with the peer, you can use one of two options:

- [Peer IP Address, page 17-27](#)—used to configure both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- [Special Frames, page 17-27](#)—used to configure one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the port 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 also 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.

To assign the peer information based on the IP address, port number, or a profile ID, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch(config-if)# peer-info ipaddr 10.1.1.1 | Assigns an IP address to configure the peer information. Since 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. |
| Step 2 | switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000 | Assigns the IP address and sets the peer TCP port to 3000. The valid port number range is from 0 to 65535. |
| | switch(config-if)# no peer-info ipaddr 10.1.1.1 port 2000# | Deletes the assigned peer port information. |
| Step 3 | switch(config-if)# peer-info profile_id 20 | Assigns the peer profile ID to connect to 20. The valid range is from 1 to 255. |
| | switch(config-if)# no peer-info profile_id 500 | Deletes the assigned peer profile information. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |

Special Frames

You can alternatively establish a FCIP link with a peer using an optional protocol called special frames. You can enable or disable the **special-frame** option. On the peer side, the **special-frame** option must be enabled in order to establish the FCIP link. When the **special-frame** option is 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 exchanged to discover and authenticate the link.

By default, the special frame feature is disabled.



Note

Refer to the Fibre Channel IP standards for further information on special frames.

Send documentation comments to mdsfeedback-doc@cisco.com

To enable special frames, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config-if)# special-frame peer-wnn 12:12:34:45:ab:bc:cd:00</code> | 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 command to obtain the peer WWN . |
| | <code>switch(config-if)# no special-frame peer-wnn 12:12:34:45:ab:bc:cd:00</code> | Disables special frames (default). |
| Step 2 | <code>switch(config-if)# special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code> | Enables special frames and sets the peer WWN as specified by the profile ID (155). |
| | <code>switch(config-if)# no special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code> | Disables special frames (default). |
| Step 3 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Configuring Active Connection

Use the **passive-mode** option to configure the required mode for initiating an IP 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 and merely 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 will not be initiated.

To enable the passive mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch(config-if)# passive-mode</code> | Enable passive mode while attempting a TCP connection. |
| | <code>switch(config-if)# no passive-mode</code> | Reverts to the factory set default of using the active mode while attempting the TCP connection. |
| Step 2 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring the Number of TCP Connections

Use the **tcp-connection** option to specify the number of TCP connections from a FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure 1 or 2 TCP connections.

For example, the Cisco PA-FC-1G Fibre Channel port adapter which has only 1 (one) TCP connection interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit and you can change the configuration on the switch using the **tcp-connection 1** command. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, the software handles it gracefully and moves on 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. Two (2) is the default and the maximum number of TCP connection attempts. |
| | 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. |

Enabling Time Stamps

Use the **time-stamp** option to enable or disable FCIP time stamps on a packet. The **time stamp** option instructs the switch to discard packets that are outside the specified time. By default, the **time-stamp** option is disabled.

The **acceptable-diff** option 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 if a packet arrives within a 1000 millisecond interval (+ or -1000 milliseconds), that packet is accepted.



Note

If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the “[NTP Configuration](#)” section on page 3-18).

To enable or disable the **time-stamp** option, follow these steps:

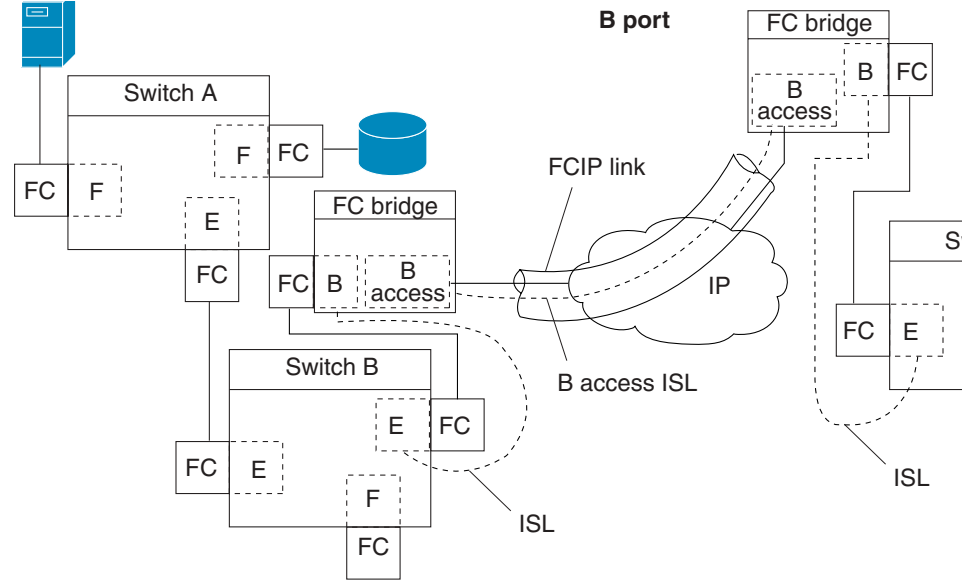
| | Command | Purpose |
|--------|---|--|
| Step 1 | switch(config-if)# time-stamp Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link | Enables time stamp checking for received packets with a default acceptable time difference of 1000 milliseconds. |
| | switch(config-if)# no time-stamp | Disables (default) time stamps. |
| Step 2 | switch(config-if)# time-stamp acceptable-diff 4000 | Configures the acceptable time within which a packet is accepted. The default difference is a 1000 millisecond interval from the network time. The valid range is from 1 to 60,000 milliseconds. |
| | switch(config-if)# no time-stamp acceptable-diff 500 | Deletes the configured time difference and reverts the difference to factory defaults. |
| Step 3 | switch(config-if)# no shutdown | Enables the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com

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 17-11](#) depicts a typical SAN extension over an IP network.

Figure 17-11 FCIP B Port and Fibre Channel E Port



B ports bridge Fibre Channel traffic from one E port to a remote E port without participating in fabric-related activities such as principal switch election, Domain ID assignment, and Fibre Channel routing (FSPF). For example, Class F traffic entering a SAN extender does not interact with the B port.

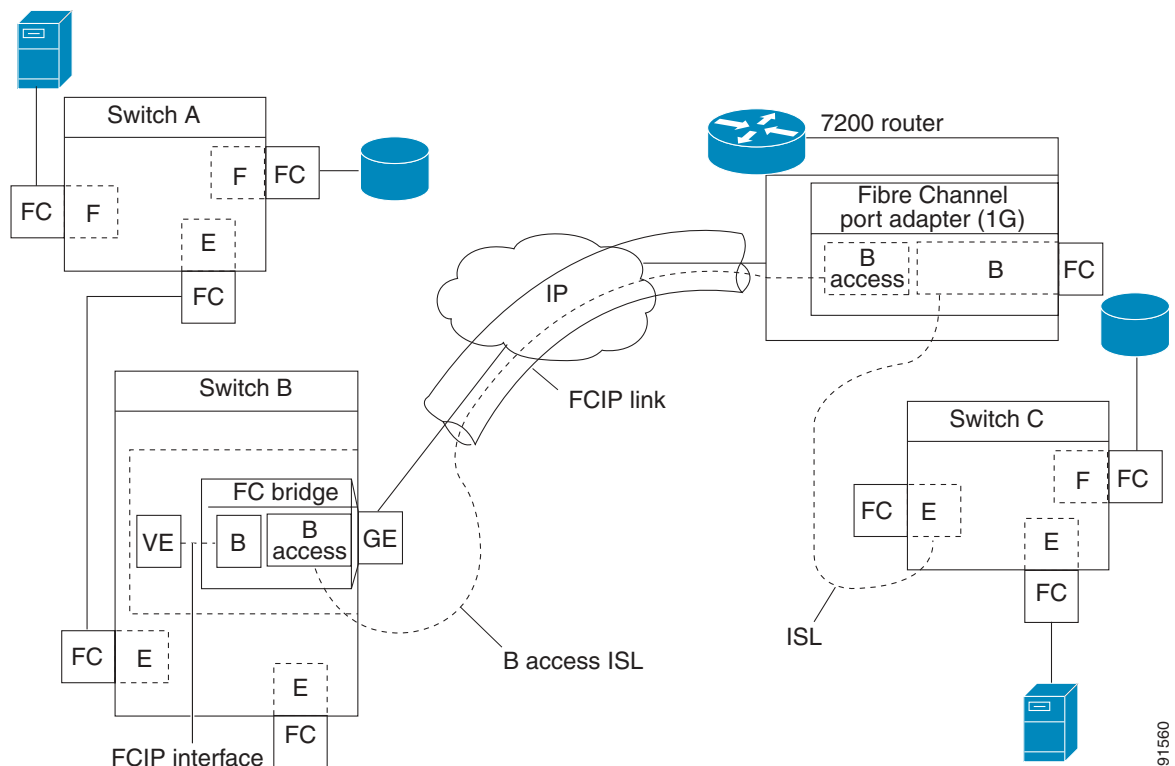
Send documentation comments to mdsfeedback-doc@cisco.com

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 which 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 a FCIP link, B ports use a B access ISL*.

The IPS module supports 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 corresponding virtual B port connects to an virtual E port which completes the end-to-end E port connectivity requirement (see Figure 17-12).

Figure 17-12 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module allows remote B port SAN extenders to communicate directly with an Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring B Ports

When a 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 | <code>switch(config-if)# bport</code> | Enables B port mode on the FCIP interface. |
| | <code>switch(config-if)# no bport</code> | Reverts to E port mode on the FCIP interface (default). |
| Step 2 | <code>switch(config-if)# bport-keepalive</code> | Enables the reception of keepalive responses sent by a remote peer. |
| Step 3 | <code>switch(config-if)# no bport-keepalive</code> | Disables the reception of keepalive responses sent by a remote peer (default). |

E Port Configurations

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available FCIP interfaces:

- VSANs (see [Chapter 8, “Configuring and Managing VSANs”](#))
 - FCIP interfaces can be a member of any VSAN.
- Trunk mode (see [Chapter 10, “Configuring Trunking”](#))
 - Trunk mode can be configured.
 - Trunk allowed VSANs can be configured
- PortChannels (see [Chapter 11, “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 15, “Configuring Fibre Channel Routing Services and Protocols”](#))
- Fibre Channel domains (fcdomains—see [Chapter 19, “Configuring Domain Parameters”](#))
- Zone merge (see [Chapter 12, “Configuring and Managing Zones”](#))
 - Importing the zone database from the adjacent switch.
 - Exporting the zone database from the adjacent switch.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Displaying FCIP Information

Use the **show interface** commands to view the summary, counter, description, and status of the FCIP link. Use the output of these commands to verify the administration mode, the interface status, the operational mode, the related VSAN ID, and the profile used. See Examples 17-14 to 17-19.

Example 17-14 Displays the FCIP Interface

```
switch# show interface fcip 3
fcip3 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:ca:00:05:30:00:07:1e
  Peer port WWN is 20:ca:00:00:53:00:18:1e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1,10)
  Trunk vsans (operational) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (10)
  Trunk vsans (initializing) ()
  Using Profile id 3 (interface GigabitEthernet4/3)
  Peer Information
    Peer Internet address is 43.1.1.1 and port is 3225
    Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  B-port mode disabled
  TCP Connection Information
    2 Active TCP connections
      Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
      Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
    30 Attempts for active connections, 0 close of connections
  TCP Parameters
    Path MTU 1500 bytes
    Current retransmission timeout is 300 ms
    Round trip time: Smoothed 10 ms, Variance: 5
    Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
    Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
    Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
  5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
  5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
    808 frames input, 75268 bytes
      808 Class F frames input, 75268 bytes
      0 Class 2/3 frames input, 0 bytes
      0 Error frames timestamp error 0
    806 frames output, 74712 bytes
      806 Class F frames output, 74712 bytes
      0 Class 2/3 frames output, 0 bytes
      0 Error frames 0 reass frames
```

Example 17-15 Displays Detailed FCIP Interface Counter Information

```
switch# show interface fcip 3 counters
fcip3
  TCP Connection Information
    2 Active TCP connections
      Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
      Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
    30 Attempts for active connections, 0 close of connections
  TCP Parameters
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

Path MTU 1500 bytes
Current retransmission timeout is 300 ms
Round trip time: Smoothed 10 ms, Variance: 5
Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
814 frames input, 75820 bytes
  814 Class F frames input, 75820 bytes
    0 Class 2/3 frames input, 0 bytes
    0 Error frames timestamp error 0
812 frames output, 75264 bytes
  812 Class F frames output, 75264 bytes
    0 Class 2/3 frames output, 0 bytes
    0 Error frames 0 reass frames

```

Example 17-16 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 |

Example 17-17 Displays the FCIP Interface Description

```
switch# show interface fcip 51 description
```

```

FCIP51
  Sample FCIP interface

```

Example 17-18 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 17-19 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

```

Send documentation comments to mdsfeedback-doc@cisco.com

Minimum available bandwidth is 15000 kbps
Estimated round trip time is 1000 usec

FCIP High Availability

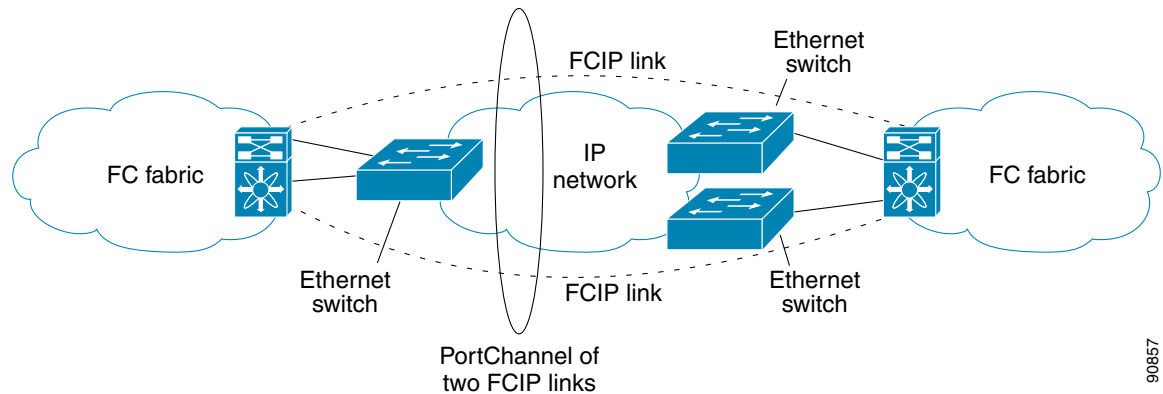
The following high availability solutions are available for FCIP configurations:

- [Fibre Channel PortChannels, page 17-35](#)
- [FSPE, page 17-36](#)
- [VRRP, page 17-36](#)
- [Ethernet PortChannels, page 17-37](#)

Fibre Channel PortChannels

[Figure 17-13](#) 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 17-13 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.


90857

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

FSPF

Figure 17-14 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 17-14 FSPF-Based Load Balancing



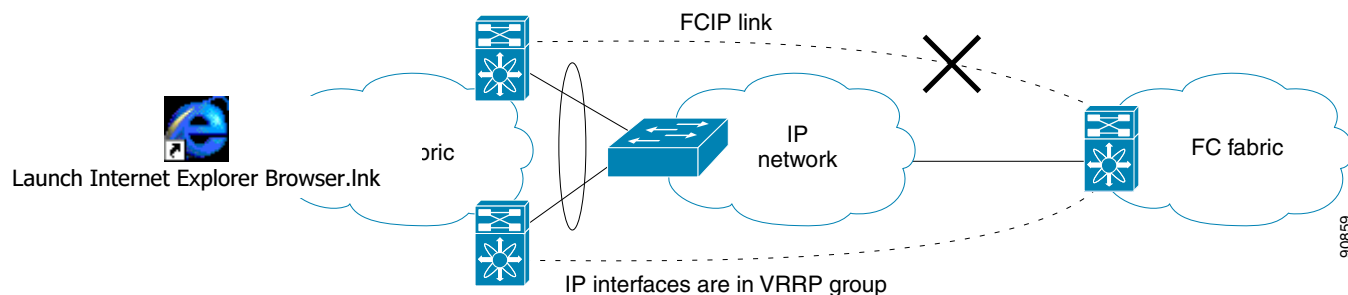
The following characteristics set FSPF solutions apart from other solutions:

- 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 17-15 displays a VRRP-based high availability FCIP configuration 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 17-15 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Ethernet PortChannels

Figure 17-16 displays a Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 17-16 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 appears like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer Ethernet-level redundancy, Fibre Channel PortChannels offer (E)ISL-level redundancy. FCIP is unaware of any Ethernet PortChannels or Fibre Channel PortChannels. Fibre Channel PortChannels are unaware of any Ethernet PortChannels, and there is no mapping between the two (see [PortChannels at the Fibre Channel and Ethernet Levels](#), page 17-37).

Figure 17-17 PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see [Chapter 11, “Configuring PortChannels.”](#) To configure Ethernet PortChannels, refer to the [“Configuring Ethernet PortChannels”](#) section on page 17-14.

Send documentation comments to mdsfeedback-doc@cisco.com

Configuring iSCSI


This section includes the following topics:

- [About iSCSI, page 17-38](#)
- [Presenting Fibre Channel Targets as iSCSI Targets, page 17-41](#)
- [iSCSI Virtual Target Configuration Examples, page 17-44](#)
- [Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 17-46](#)
- [Access Control in iSCSI, page 17-49](#)
- [User Authentication Using iSCSI, page 17-51](#)
- [Assigning VSAN Membership to iSCSI Hosts, page 17-48](#)
- [Displaying iSCSI Information, page 17-53](#)
- [iSCSI High Availability, page 17-62](#)
- [iSCSI Authentication Setup Guidelines, page 17-64](#)

About iSCSI

The IPS module provides transparent SCSI routing. IP hosts using iSCSI protocol can transparently access iSCSI targets on the Fibre Channel network. [Figure 17-18](#) provides an example of a typical configuration of iSCSI hosts with access to a Fibre Channel SAN.


Figure 17-18 Typical IP to Fibre Channel SAN Configuration



Send documentation comments to mdsfeedback-doc@cisco.com


The IPS module enables you to create virtual iSCSI targets and maps them to physical Fibre Channel targets available in the Fibre Channel SAN. It presents the Fibre Channel targets to IP hosts as if the physical targets were attached to the IP network (see Figure 17-19).

Figure 17-19 iSCSI View



In conjunction with presenting Fibre Channel targets to iSCSI hosts, the iSCSI feature presents each iSCSI host as a Fibre Channel host, i.e. Host Bus Adaptor (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network (see Figure 17-20).

Figure 17-20 Fibre Channel SAN View



Note

Refer to the IETF standards for IP storage at <http://www.ietf.org>, for information on the iSCSI protocol.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Routing iSCSI Requests and Responses

The iSCSI feature consists of routing iSCSI requests and responses between 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 17-21](#)).

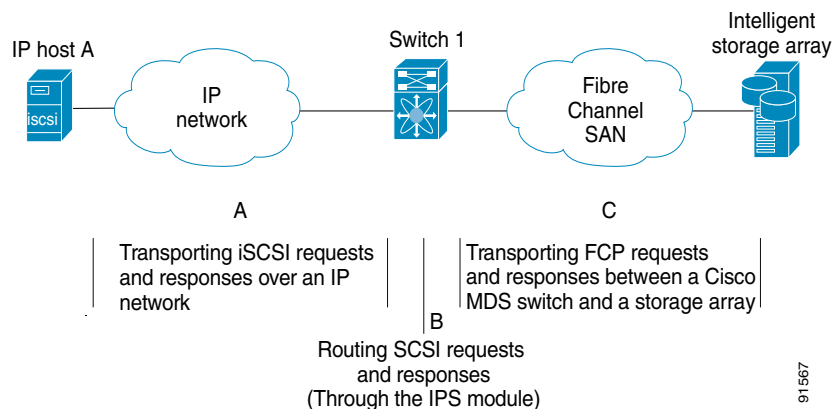
Figure 17-21 Routing iSCSI Requests and Responses for Transparent iSCSI Routing

Each iSCSI host that requires access to storage via the IPS module needs to have a compatible iSCSI driver installed. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the perspective of a host operating system, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver for a peripheral channel in the host. From the perspective of the storage device, each IP host appears as a Fibre Channel host.

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions (see [Figure 17-21](#)):

- Transporting iSCSI requests and responses over an IP network between hosts and the IPS module.
- Routing SCSI requests and responses between hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). This routing is performed by the IPS module.
- Transporting FCP requests or responses between the IPS module and Fibre Channel storage devices.

Figure 17-22 Transparent SCSI Routing Actions



Note

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN.

Send documentation comments to mdsfeedback-doc@cisco.com

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module presents physical Fibre Channel targets as iSCSI targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- **Dynamic Importing**—used if all logical units (LUs) in all Fibre Channel storage targets are made available to iSCSI hosts (subject to VSAN and zoning).
- **Static Importing**—used if iSCSI hosts are restricted to subsets of LUs in the Fibre Channel targets and additional iSCSI access control is needed (see the “[Access Control in iSCSI](#)” section on [page 17-49](#)). Also, static import allows automatic failover if the Fibre Channel targets’ LU is reached by redundant Fibre Channel ports (see the “[High Availability Static Importing](#)” section on [page 17-43](#)).



Note

The IPS module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have a configured name. Un mapped targets are advertised with the name created by the conventions explained in this section.

Dynamic Importing

To enable dynamic importing of Fibre Channel targets into iSCSI, use the **iscsi import target fc** command.

The IPS module maps each physical Fibre Channel target port as one iSCSI target. That is, all LU accessible via the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the storage target.

For example, if an iSCSI target was created for Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0 through 2, those LUNs would become available to an IP host as LUNs 0 through 2 as well.

The iSCSI target node name is created automatically using the iSCSI I qualified name (IQN) format. The IPS module creates an IQN formatted iSCSI node name using the following conventions:

- IPS ports that are not part of a VRRP group 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.1987-05.com.cisco:05.PC-<port-ch-intf#>-<port-ch-sub-intf#>.<Target-pWWN>
```



Note

In this format, each IPS port in a Cisco MDS 9000 Family switch creates a different iSCSI target node name for the same Fibre Channel target.

Send documentation comments to mdsfeedback-doc@cisco.com

To dynamically import Fibre Channel targets, 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 dynamically map each Fibre Channel target in the Fibre Channel SAN to the IP network. The automatically-created iSCSI target node names use the IQN format. Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms. |

Static Importing

You can manually (statically) create an iSCSI target and assign a node name to it. A statically-mapped iSCSI target can either contain the whole FC target port, or it can contain one or more LUs from a Fibre Channel target port.

To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name iqn.abc switch(config-(iscsi-tgt))# | Creates the iSCSI target name iqn.abc. |
| Step 3 | switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06 | Maps a virtual target node to a Fibre Channel target. One iSCSI target cannot contain more than one Fibre Channel target. Don't specify the LUN if you wish to map the whole Fibre Channel target to an iSCSI target. All Fibre Channel target LUNs are exposed to iSCSI. Use the LUN option to map different Fibre Channel LUNs to different iSCSI virtual targets. If you have already mapped the whole Fibre Channel target, you will not be able to use this option. |
| | switch(config-(iscsi-tgt))# pwwn 26:00:00:00:00:11:00:11 fc-lun 1 iscsi-lun 1 | Maps the whole target using LUN mapping options. |

Creating iSCSI Targets

To create a static iSCSI target for the entire Fibre Channel target port, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name iqn.abc switch(config-(iscsi-tgt))# | Creates the iSCSI target name iqn.abc. |

Send documentation comments to mdsfeedback-doc@cisco.com

Advertising iSCSI Targets

You can limit the Gigabit Ethernet interfaces over which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config-(iscsi-tgt))# advertise interface GigabitEthernet 2/5</code> | Advertises the virtual target only on the specified interface. By default, it is advertised on all interfaces in all IPS modules. |
| | <code>switch(config-(iscsi-tgt))# no advertise interface GigabitEthernet 2/5</code> | Removes this interfaces from the list of interfaces from which this target is advertised. |

High Availability Static Importing

Physical Fibre Channel targets are configured to have LUs visible over two Fibre Channel N ports— one in active mode and another in passive mode. When the active port fails, the passive port takes over. 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 passive port becomes active and the iSCSI session switches to use the new active port. If both the primary and secondary pWWNs are available, then both pWWNs can be used—each session may use either pWWN (see [Figure 17-23](#)).

Figure 17-23 Mapping LUNs to be Visible

In [Figure 17-23](#), you can create a virtual iSCSI target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

To create a static iSCSI virtual target, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# iscsi virtual-target name abc</code> | Creates the iSCSI target name iqn.abc. |
| Step 3 | <code>switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06</code> <code>secondary-pwwn 26:00:01:02:03:10:11:12</code> | Configures the secondary port for this virtual target. |

Send documentation comments to mdsfeedback-doc@cisco.com

iSCSI Virtual Target Configuration Examples


This section provides three examples of virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a virtual iSCSI target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 17-24](#)).

```
iscsi virtual-target name iqn.abc-1
pWWN 28:00:01:02:03:04:05:06
```

Figure 17-24 Assigning iSCSI Node Names




91569

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 17-25](#)).

```
iscsi virtual-target name iqn.abc-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 1
iscsi virtual-target name iqn.abc-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 1
iscsi virtual-target name iqn.abc-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 1
```

Figure 17-25 Mapping LUNs to iSCSI a Node Name



91570


Send documentation comments to mdsfeedback-doc@cisco.com

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 17-26](#)).

```
iscsi virtual-target name iqn.abc-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.abc-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.abc-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
```

Figure 17-26 Mapping LUNs to Multiple iSCSI Node Names



91571

Send documentation comments to mdsfeedback-doc@cisco.com

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The iSCSI hosts are mapped to virtual Fibre Channel hosts in one of two ways (see [Figure 17-20](#)):

- **Dynamic Mapping** (default)—used if no access control is done on the Fibre Channel target. An iSCSI host may use different pWWNs each time it connects to a Fibre Channel target.
- **Static Mapping**—used if an iSCSI host should always have the same pWWN or nWWN each time it connects to a Fibre Channel target.

Dynamic Mapping

When an iSCSI host connects to the IPS module using the iSCSI protocol, a virtual N port is created for the host. The nWWNs and pWWNs are dynamically allocated from the switch's Fibre Channel WWN pool. The IPS module registers this N port in the Fibre Channel SAN. The IPS module continues using that nWWN and pWWN to represent this iSCSI host until it no longer has a connection to any iSCSI target via that IP storage port.

At that point, the virtual Fibre Channel host is taken offline from the Fibre Channel SAN and the nWWNs and pWWNs are released back to the switch's Fibre Channel WWN pool. These addresses become available for assignment to other iSCSI hosts requiring access to Fibre Channel SANs.

When a dynamically mapped iSCSI initiator has multiple sessions to multiple Fibre Channel targets, each session can use the same pWWN and nWWN as long as it uses the same node name in the iSCSI login message. If the host has multiple network interfaces (and the same IP address), and each IP address is treated as different hosts, then the **switchport initiator id ip-address** command is used to identify an iSCSI initiator. This command uses the IP address instead of the initiator name.

All dynamic iSCSI initiators are members of the default VSAN (VSAN 1).

Identifying Initiators

An iSCSI initiator is identified in one of two ways:

- The iSCSI node name (**switchport initiator id name** command)—If the node name is used, an initiator with multiple IP addresses (multiple interface cards—NICs or multiple network interfaces) has one virtual N port.
- The IP address (**switchport initiator id ip-address** command)—If the IP address is used, a virtual N port is created for each NIC or network interface.

By default, the switch uses the iSCSI node name to identify the initiator. You can change this default so the switch identifies the initiator using the IP address.

To identify the initiator using the IP address, 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. |

Send documentation comments to mdsfeedback-doc@cisco.com

Static Mapping

With dynamic mapping, each time the iSCSI host connects to the IPS module a new Fibre Channel N port is created and the nWWNs and pWWNs allocated for this N port may be different. Use the static mapping method if you need to obtain the same nWWNs and pWWNs for the iSCSI host each time it connects to the IPS module.

You can implement static mapping in one of two ways: system assignment or manual assignment.

- System assignment—When a static mapping configuration is created, one nWWN and/or one or more pWWNs are allocated from the switch's Fibre Channel WWN pool and the mapping is kept permanent. This assignment uses the **system-assign** option.
- Manual assignment—You can specify your own unique WWN using the **manual-assign** option. Each time the iSCSI session is created, the same nWWN/pWWN that was initially created is used.



Tip

We recommend using the **system-assign** option. If you manually assign a WWN, you must uniquely associate the WWN to a single device (see the [“Configuring World Wide Names”](#) section on page 24-16).

Static mapping can be used on the IPS module to access intelligent Fibre Channel storage arrays that have access control and LUN mapping/masking configuration based on the initiator's pWWNs and/or nWWNs.



Note

If an iSCSI host connects to multiple IPS ports, each port independently creates one virtual N port for the host. If static mapping is used, enough pWWNs should be configured for as many IPS ports to which a host connects.

To configure static mapping (using the **name** 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 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 255 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 IP address of the initiator node. |
| | switch(config)# no iscsi initiator ip address 10.50.0.0 | Deletes the configured iSCSI initiator. |

Send documentation comments to mdsfeedback-doc@cisco.com

To assign the WWN for an iSCSI initiator, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-(iscsi-init))# static nWWN system-assign</code> | Uses the switch's WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent. |
| | <code>switch(config-(iscsi-init))# nWWN 20:00:00:05:30:00:59:11</code> | Assigns the user provided WWN as nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node. |
| Step 2 | <code>switch(config-(iscsi-init))# static pWWN system-assign 2</code> | Uses the switch's WWN pool to allocate two pWWNs for this iSCSI initiator and keeps it persistent. The range is from 1 to 64. |
| | <code>switch(config-(iscsi-init))# pWWN 21:00:00:20:37:73:3b:20</code> | Assigns the user provided WWN as pWWN for the iSCSI initiator. |

Assigning VSAN Membership to iSCSI Hosts

An iSCSI host can reside in multiple VSANs based on the configuration. By default, a host is only in VSAN 1 (default VSAN). The IPS module creates one Fibre Channel virtual N port in each VSAN to which the host belongs.

An iSCSI host can become a member of one or more VSANs.

To assign VSAN membership for iSCSI hosts, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator</code> <code>switch(config-(iscsi-init))#</code> | Configures an iSCSI initiator. |
| Step 3 | <code>switch(config-(iscsi-init))# vsan 3</code> | Assigns the iSCSI initiator node to a specified VSAN. Note You can assign this host to one or more VSANs. |
| | <code>switch(config-(iscsi-init))# no vsan 5</code> | Removes the iSCSI node from the specified VSAN. |



Note

By default, an iSCSI initiator is only present in the default VSAN (VSAN 1). 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Access Control in iSCSI

You can control access to each statically-mapped iSCSI target by specifying a list of IPS ports on which it will be advertised and specifying a list of iSCSI initiator node names allowed to access it. Fibre Channel zoning-based access control and iSCSI-based access control are the two mechanisms by which access control can be provided for iSCSI. Both methods can be used simultaneously.



Note

This access control is in addition to the existing Fibre Channel access control. The iSCSI initiator has to be in the same VSAN and zone as the physical Fibre Channel target.

Fibre Channel Zoning-Based Access Control

Zoning is an access control mechanism within a VSAN. The switch's zoning implementation extends the VSAN and zoning concepts from the Fibre Channel domain to also cover the iSCSI domain. This extension includes both iSCSI and Fibre Channel features and provides a uniform, flexible access control across a SAN. Static and dynamic are the two Fibre Channel zoning access control mechanisms.

- **Static**—statically map the iSCSI host to Fibre Channel virtual N port(s). This creates a permanent nWWNs and pWWNs. Next, configure the assigned pWWN into zones, similar to adding a regular Fibre Channel host's pWWN to a zone.
- **Dynamic**—add the iSCSI host's initiator node name as a member of a zone. When the IP host's Fibre Channel virtual N port is created and the Fibre Channel address (nWWNs and pWWNs) is assigned, Fibre Channel zoning is enforced.

To register an iSCSI initiator in 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 to be included. |
| Step 3 | switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1 | Adds the device as specified by the node name. |
| | switch(config-zone)# no member iqn.1987-02.com.cisco.initiator1 | Deletes the specified device. |
| | switch(config-zone)# member symbolic-nodename 10.50.1.1 | Adds the device as specified by the IP address. |
| | switch(config-zone)# no member 10.50.1.1 | Deletes the specified device. |

iSCSI-Based Access Control

For static iSCSI targets, you can manually configure a list of iSCSI initiators that are allowed to access it. The iSCSI initiator is identified by the iSCSI node name or the IP address of the iSCSI host.

By default, static virtual iSCSI targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a virtual iSCSI target to be accessed by all hosts. The initiator access list can contain one or more initiators. Each initiator is identified by one of the following:

- iSCSI node names
- IP addresses
- IP subnets

Send documentation comments to mdsfeedback-doc@cisco.com

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.abc switch(config-(iscsi-tgt))# | Creates the iSCSI target name iqn.abc. |
| 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 IP 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 IP address from accessing virtual targets. |
| | switch(config-(iscsi-tgt))# initiator ip address 10.50.1.1 255.255.255.0 permit | Allows all initiators in this subnetwork to access this virtual target. |
| | switch(config-(iscsi-tgt))# no initiator ip address 10.50.1.1 255.255.255.0 permit | Prevents all initiators in this 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 use both iSCSI node name-based and Fibre Channel zoning-based access control lists to enforce access control during iSCSI discovery and iSCSI session creation.

- iSCSI discovery—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS 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.
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module verifies if the specified iSCSI target (in the session login request) is a static mapped target, and if true, verifies if the IP host's iSCSI node name is allowed to access the target. If the IP host does not have access, its login is rejected.

The IPS module, then creates a Fibre Channel virtual N port (the N port may already exist) for this IP host and does a Fibre Channel name server query for the FCID of the Fibre Channel target pWWN that is being accessed by the IP host. It uses the IP host virtual N port's pWWN as the requester of the name server query. Thus, the name server does a zone-enforced query for the pWWN and responds to the query.

If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

User Authentication Using iSCSI

The IPS module supports the iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. When iSCSI authentication is enabled, the iSCSI hosts must provide user name and password information each time an iSCSI session is established.



Note

Only the Challenge Handshake Authentication Protocol (CHAP) authentication method is supported.

The IPS module also supports iSCSI hosts to challenge the IPS module to authenticate itself.

The **aaa auth iscsi radius** command enables RADIUS authentication for the iSCSI host. If RADIUS authentication is not enabled or RADIUS servers are unavailable, the local database is used (see the “Configuring RADIUS Authentication” section on page 14-14).

To configure RADIUS 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 radius | Uses RADIUS for iSCSI authentication method. |
| | switch(config)# aaa authentication iscsi local | Configures the switch to only use the local password database for iSCSI CHAP authentication. |

Authentication Mechanism

During an iSCSI login, both the iSCSI initiator and target have the option to authenticate each other. By default, the IPS module allows either CHAP authentication or no authentication from iSCSI hosts. If CHAP authentication should always be used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used, issue the **iscsi authentication none** command.



Note

The authentication for a Gigabit Ethernet interface or subinterface configuration overrides the authentication for the global interface configuration.

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. The validation is done using RADIUS or local authentication. |

Send documentation comments to mdsfeedback-doc@cisco.com

To configure the authentication policy 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. |

The IPS module verifies the iSCSI host authentication in one of two ways: the local password database or RADIUS (see the “[User Authentication](#)” section on page 14-2). If local authentication is used, the **username iscsi-user password iscsi** command assigns a password and a user name for a new user. If the user name does not exist it will be created.



Note

The **iscsi** keyword is mandatory to identify iSCSI users.

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 ffsffsfssfs345353554535 iscsi | Configures a user name (iscsiuser) and password (ffsffsfssfs345353554535) in the local database for iSCSI login authentication. |
| Note | The iscsi keyword is required at the end to identify the user. | |

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Displaying iSCSI Information

This section includes the following topics:

- [Displaying iSCSI Interfaces, page 17-53](#)
- [Displaying Global iSCSI Information, page 17-55](#)
- [Displaying iSCSI Sessions, page 17-55](#)
- [Displaying iSCSI Initiators, page 17-56](#)
- [Displaying iSCSI Virtual Targets, page 17-59](#)
- [Displaying IPS Statistics, page 17-60](#)
- [Displaying iSCSI User Information, page 17-61](#)

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. See [Example 17-20](#).

Example 17-20 Displays the iSCSI Interface Information

```
switch# show interface iscsi 2/1
iscsi2/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:41:00:05:30:00:50:de
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 7, Number of TCP connection: 7
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is disabled
    Keepalive-timeout is 1 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 8
    Sack is disabled
    Minimum available bandwidth is 0 kbps
    Estimated round trip time is 0 usec
  5 minutes input rate 265184 bits/sec, 33148 bytes/sec, 690 frames/sec
  5 minutes output rate 375002168 bits/sec, 46875271 bytes/sec, 33833 frames/sec
  iSCSI statistics
    6202235 packets input, 299732864 bytes
      Command 6189718 pdus, Data-out 1937 pdus, 1983488 bytes, 0 fragments
    146738794 packets output, 196613551108 bytes
      Response 6184282 pdus (with sense 4), R2T 547 pdus
      Data-in 140543388 pdus, 189570075420 bytes
```

The **show iscsi stats** command can be used to view brief or detailed iSCSI statistics per iSCSI interface. See Examples [17-21](#) and [17-22](#).

Example 17-21 Displays iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 4/1
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
iSCSI statistics
  1196 packets input, 173680 bytes
    Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
  output 1802 packets, 647152 bytes
    Response 483 pdus (with sense 0), R2T 25 pdus
    Data-in 685 pdus, 554696 bytes
```

Example 17-22 Displays Detailed iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 4/1 detail
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  iSCSI statistics
    1196 packets input, 173680 bytes
      Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
    output 1802 packets, 647152 bytes
      Response 483 pdus (with sense 0), R2T 25 pdus
      Data-in 685 pdus, 554696 bytes
  iSCSI Forward:
    Command: 483 PDUs (Rcvd: 483)
    Data-Out (Write): 104 PDUs (Rcvd 104), 0 fragments, 106496 bytes
  FCP Forward:
    Xfer_rdy: 25 (Rcvd: 25)
    Data-In: 685 (Rcvd: 719), 554696 bytes
    Response: 483 (Rcvd: 534), with sense 0
    TMF Resp: 0

  iSCSI Stats:
    Login: attempt: 25, succeed: 25, fail: 0, authen fail: 0
    Rcvd: NOP-Out: 556, Sent: NOP-In: 556
      NOP-In: 0, Sent: NOP-Out: 0
      TMF-REQ: 0, Sent: TMF-RESP: 0
      Text-REQ: 6, Sent: Text-RESP: 6
      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
  FCP Stats:
    Total: Sent: 726
      Received: 1366 (Error: 0, Unknown: 0)
    Sent: PLOGI: 17, Rcvd: PLOGI_ACC: 17, PLOGI_RJT: 0
      PRLI: 17, Rcvd: PRLI_ACC: 17, PRLI_RJT: 0, Error resp: 0
      LOGO: 12, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
      PRLO: 12, Rcvd: PRLO_ACC: 0, PRLO_RJT: 0
      ABTS: 0, Rcvd: ABTS_ACC: 0
      TMF REQ: 0
      Self orig command: 51, Rcvd: data: 34, resp: 51
    Rcvd: PLOGI: 20, Sent: PLOGI_ACC: 5, PLOGI_RJT: 15
      LOGO: 5, Sent: LOGO_ACC: 5, LOGO_RJT: 0
      PRLI: 5, Sent: PRLI_ACC: 5, PRLI_RJT: 0
      PRLO: 0, Sent: PRLO_ACC: 0, PRLO_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
      Persistent Resv 0, No task: 0
    Data-Out: 0, Data CRC Error: 0
    TMF-Req: 0, No task: 0
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
FCP Drop:
  Xfer_rdy: 0, Data-In: 0, Response: 0

Buffer Stats:
  Buffer less than header size: 0, Partial: 53, Split: 79
  Pullup give new buf: 0, Out of contiguous buf: 0, Unaligned m_data: 0
```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See [Example 17-23](#)

Example 17-23 Displays the Current Global iSCSI Configuration and State.

```
switch# show iscsi global
iSCSI Global information
  Authentication: NONE
  Import FC Target: Enabled
  Number of target nodes: 5
  Number of portals: 8
  Number of sessions: 6
  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 specificity an initiator, a target, or both.

[Example 17-24](#) displays one iSCSI initiator configured based on the iqn name (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on its IP address (10.10.100.199).

Example 17-24 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
Session #3
Target iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
VSAN 1, ISID 246e00000000, Status active, no reservation
```

[Example 17-25](#) and [Example 17-26](#) display the iSCSI initiator configured based on its IP address (10.10.100.199).

Example 17-25 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 17-26 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 a 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 **fcplib-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See [Examples 17-27](#) and [17-28](#).

Example 17-27 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

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: ign.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 17-28 Display Detailed Information About the iSCSI Initiator

```

switch# show iscsi initiator ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is ign.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
    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
  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

```

Send documentation comments to mdsfeedback-doc@cisco.com

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 Examples 17-29 and 17-31.

Example 17-29 Displays Fibre Channel Name Server Database

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x6c0001      NL    21:00:00:04:cf:4c:52:c1 (Seagate)          scsi-fcp:target
0x6c01e8      NL    21:00:00:20:37:62:c0:0c (Seagate)          scsi-fcp:target
0x6c0202      N     22:04:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
0x6c0203      N     22:00:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
0x6c0301      NL    21:00:00:20:37:a6:be:32 (Seagate)          scsi-fcp:target
Total number of entries = 5

VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x6e0000      N     22:04:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
Total number of entries = 1

VSAN 5:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x640000      N     22:00:00:05:30:00:10:e1 (Cisco)            scsi-fcp:init isc..w
Total number of entries = 1
```

Example 17-30 Displays Detailed Information for a Fibre Channel N Port Created for An iSCSI Initiator Identified by it's IQN Name

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
node-wwn               :22:03:00:05:30:00:10:e1
class                  :2,3
node-ip-addr           :10.10.100.199
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsgw
symbolic-port-name     :
symbolic-node-name     :10.10.100.199
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :20:c1:00:05:30:00:10:de
hard-addr              :0x000000

Total number of entries = 1
```

Example 17-31 Displays Detailed Information for a Fibre Channel N Port created for An iSCSI Initiator Identified by it's IP Address

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
node-wwn          :22:03:00:05:30:00:10:e1
class             :2,3
node-ip-addr      :10.10.100.199
ipa               :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :10.10.100.199
port-type         :N
port-ip-addr      :0.0.0.0
fabric-port-wwn   :20:c1:00:05:30:00:10: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 17-32](#).

Example 17-32 Display 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
```

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the FC targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See [Example 17-33](#).

Example 17-33 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

Displaying IPS Statistics

The **show ips stats tcp interface** command displays information about the underlying transport for iSCSI. See Examples 17-34 and 17-35.

Example 17-34 Displays iSCSI Stats (brief)

```
switch# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
  Connection Stats
    0 active openings, 6 accepts
    0 failed attempts, 0 reset received, 6 established
  Segment stats
    640780835 received, 150953931 sent, 12 retransmitted
    0 bad segments received, 0 reset sent
  TCP Active Connections

```

| Local Address | Remote Address | State | Send-Q | Recv-Q |
|-------------------|-------------------|-----------|--------|--------|
| 10.48.69.250:3260 | 10.48.69.226:1026 | ESTABLISH | 0 | 0 |
| 10.48.69.250:3260 | 10.48.69.231:1026 | ESTABLISH | 0 | 0 |
| 10.48.69.250:3260 | 10.48.69.231:1033 | ESTABLISH | 0 | 0 |
| 10.48.69.250:3260 | 10.48.69.226:1038 | ESTABLISH | 0 | 0 |
| 0.0.0.0:3260 | 0.0.0.0:0 | LISTEN | 0 | 0 |

Example 17-35 Displays SCSI Stats (detail)

```
switch# show ips stats tcp interface gigabitethernet 2/1 detail
TCP Statistics for port GigabitEthernet2/1
  TCP send stats
    150953931 segments, 2755572300 bytes
    53986369 data, 82341597 ack only packets
    4 control (SYN/FIN/RST), 0 probes, 14625949 window updates
    12 segments retransmitted, 576 bytes
    12 retransmitted while on ethernet send queue, 0 packets split
    118741734 delayed acks sent
  TCP receive stats
    640780835 segments, 640325552 data packets in sequence, 925034009772 bytes in
sequence
    0 predicted ack, 615117910 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, 0 out-of-order packets
    0 packet after window, 0 bytes after window
    0 packets after close
    25656078 acks, 2755572210 ack bytes, 0 ack toomuch, 5786 duplicate acks
    0 ack packets left of snd_una, 0 non-4 byte aligned packets
    12100 window updates, 0 window probe
    29 pcb hash miss, 17 no port, 0 bad SYN, 0 paws drops
  TCP Connection Stats
    0 attempts, 6 accepts, 6 established
    4 closed, 4 drops, 0 conn drops
    0 drop in retransmit timeout, 4 drop in keepalive timeout
    0 drop in persist drops, 0 connections drained
  TCP Miscellaneous Stats
    21635776 segments timed, 21642712 rtt updated
    12 retransmit timeout, 0 persist timeout
    8494 keepalive timeout, 8490 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
```


Send documentation comments to mdsfeedback-doc@cisco.com

```
TCP SYN Cache Stats
  6 entries, 6 connections completed, 0 entries timed out
  0 dropped due to overflow, 0 dropped due to RST
  0 dropped due to ICMP unreachable, 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
  10.48.69.250:3260   10.48.69.226:1026   ESTABLISH   0        0
  10.48.69.250:3260   10.48.69.231:1026   ESTABLISH   0        0
  10.48.69.250:3260   10.48.69.231:1033   ESTABLISH   0        0
  10.48.69.250:3260   10.48.69.226:1038   ESTABLISH   0        0
  0.0.0.0:3260        0.0.0.0:0           LISTEN      0        0
```

The **show ips stats buffer** command displays information about the iSCSI buffers. See Example 17-36

Example 17-36 Displays iSCSI Buffers

```
switch# show ips stats buffer interface gigabitethernet 4/2
Mbuf Statistics for port GigabitEthernet4/2
Free Mbufs           : 83221
Mbuf high watermark   : 124830
Mbuf low watermark    : 20805
Mbuf alloc failures   : 0
Total clusters        : 2304
Free Clusters         : 80145
Clusters high watermark : 87381
Clusters low watermark : 79059
Clusters alloc failures : 0
Free shared mbufs     : 0
Shared Mbuf alloc failures : 0
Free shared clusters  : 0
Shared clusters alloc failures: 0

Ether channel Statistics for port GigabitEthernet4/2
TCP segments sent      : 0
TCP segments received  : 0
Xmit packets sent      : 0
Xmit packets received  : 0
Config packets sent    : 0
Config packets received : 0
MPQ packets send errors : 0
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See Example 17-37.

Example 17-37 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdfg

username:user2
secret:cshadhdhsadadjajdjas
```

Send documentation comments to mdsfeedback-doc@cisco.com

iSCSI High Availability

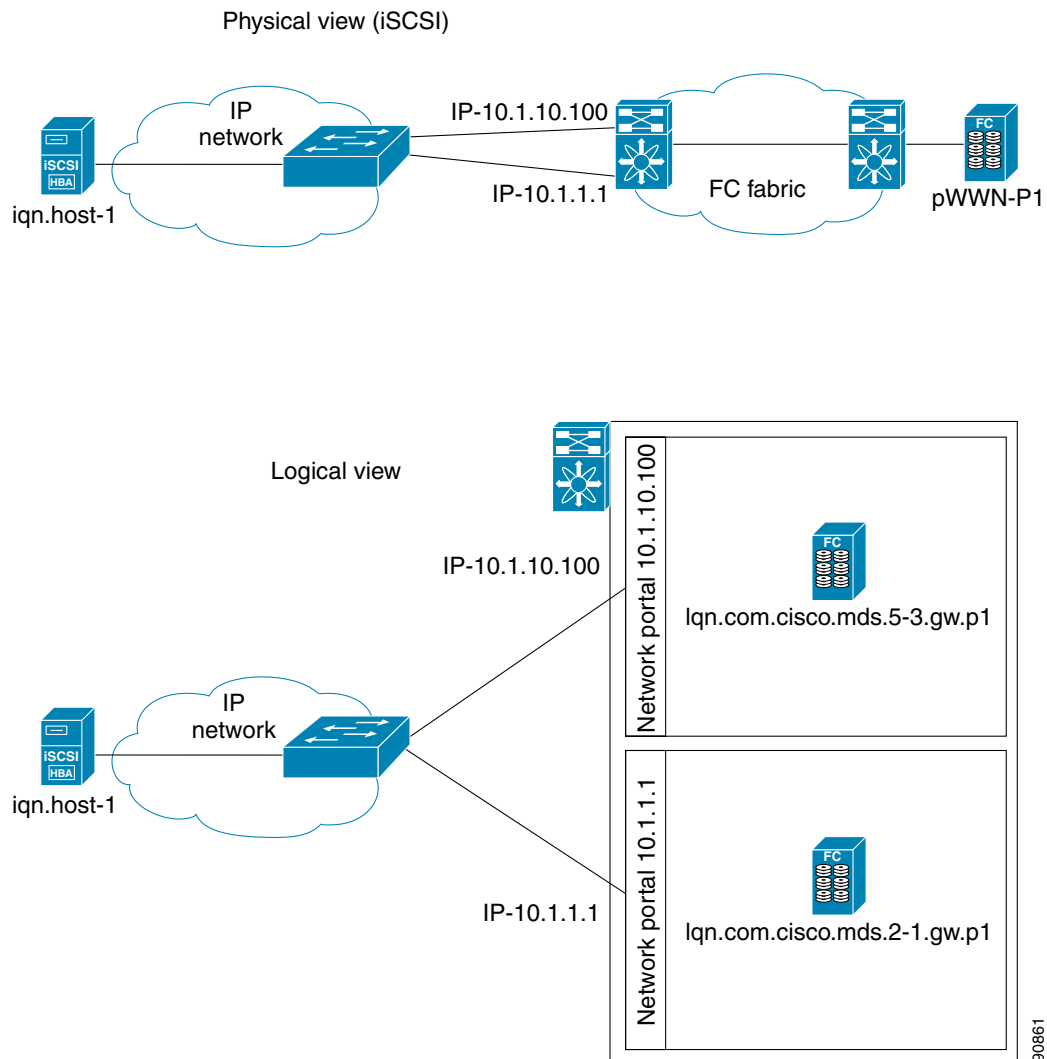
The following high availability features are available for iSCSI configurations:

- [Multiple IPS Ports Connected to the Same IP Network](#), page 17-62
- [VRRP-Based High Availability](#), page 17-63
- [Ethernet PortChannel-Based High Availability](#), page 17-64

Multiple IPS Ports Connected to the Same IP Network

Figure 17-27 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 17-27 Multiple Gigabit Ethernet Interfaces in the Same IP Network



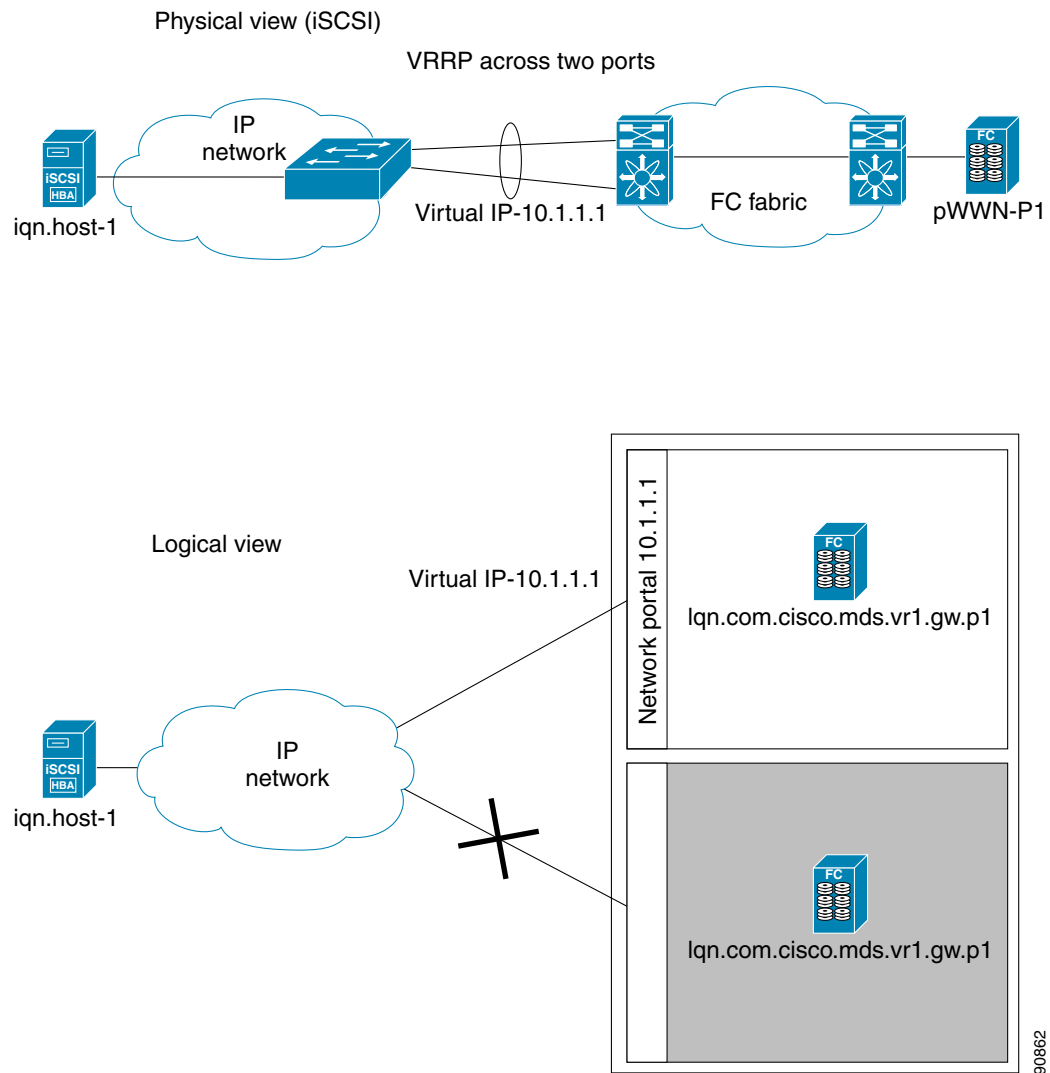
Send documentation comments to mdsfeedback-doc@cisco.com

In Figure 17-27, 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 17-28 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 17-28 VRRP-Based iSCSI High Availability



In Figure 17-28, 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Ethernet PortChannel-Based High Availability




Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that iSCSI link.

Figure 17-29 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 17-29 Ethernet PortChannel-Based iSCSI High Availability



In Figure 17-29, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the virtual iSCSI 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.

iSCSI Authentication Setup Guidelines

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- [No Authentication, page 17-64](#)
- [CHAP with Local Password Database, page 17-65](#)
- [CHAP with External RADIUS Server, page 17-65](#)
- [Scenario 1, page 17-66](#)
- [Scenario 2, page 17-71](#)



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.

No Authentication

To configure a network with no authentication set the iSCSI authentication method to none.

```
switch(config)# iscsi authentication none
```

Send documentation comments to mdsfeedback-doc@cisco.com

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 iSCSI protocol.

```
switch(config)# aaa authentication iscsi 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 MDS switch login 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** Setup the authentication verification for iSCSI protocol to go to RADIUS server.

```
switch(config)# aaa authentication iscsi radius
```

- Step 2** Configure the RADIUS server IP address.

```
switch(config)# radius-server host 10.1.1.10
```

- Step 3** Password for the MDS as RADIUS client to the RADIUS server.

```
switch(config)# radius-server key mds-1
```

- Step 4** Setup 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 CHAP.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----- Verify
  Import FC Target: Disabled
  ...
```

- Step 6** Verify that the RADIUS shared secret is MDS-1.

```
switch# show radius-server
Global RADIUS shared secret:mds-1 <----- Verify
  ...
```

Send documentation comments to mdsfeedback-doc@cisco.com

To configure an iSCSI RADIUS server, follow these steps:

- Step 1** Configure the RADIUS server to allow access from the MDS switch's management Ethernet IP address.
- Step 2** Configure the shared secret for the RADIUS server to authenticate the MDS switch.
- Step 3** Configure the iSCSI users and passwords on the RADIUS server.
- Step 4** Verify that the switch can communicate using the switch's management Ethernet IP address.

```
switch# show iscsi global
global iscsi authentication method is CHAP <----- Authentication is CHAP
...
```

- Step 5** Verify that the RADIUS server can communicate using the switch's management Ethernet IP address.

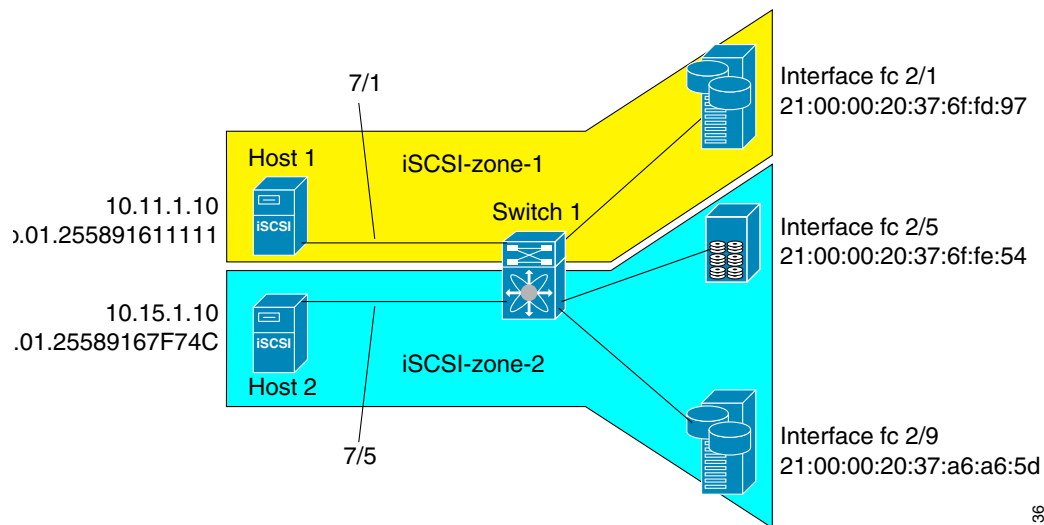
```
switch# show radius-server
Global RADIUS shared secret:mds-1 <----- Secret is mds-1
retransmission count:1
timeout value:1
following RADIUS servers are configured:
    10.1.1.100: <----- RADIUS server's IP address
...
```

Scenario 1

Sample scenario 1 assumes the following configuration (see [Figure 17-30](#)):

- Access control using Fibre Channel zoning.
- No target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator identified using IP address (Host 1 = 10.11.1.10 and Host 2 = 10.15.1.11).
- iSCSI initiator identified using node name (Host 1 = iqn.1987-05.com.cisco:01.e41695d16b1a and Host 2 = iqn.1987-05.com.cisco:01.25589167f74c).

Figure 17-30 iSCSI Scenario 1



34136

Send documentation comments to mdsfeedback-doc@cisco.com

To configure scenario 1 (see [Figure 17-30](#)), 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 FC 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 IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```

- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by the IP address, and enable the interface.

```
switch(config)# int 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 IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

- 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)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

- Step 7** Verify the available Fibre Channel targets (see [Figure 17-30](#)).

```
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 pwn 21:00:00:20:37:6f:fd:97
switch(config-zone)# member symbolic-nodename 10.11.1.10
```

Send documentation comments to mdsfeedback-doc@cisco.com

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two FC 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 zoneset 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 zoneset.

```
switch(config)# zoneset activate name zoneset-iscsi vsan 1
```

Step 12 Display the active zoneset.



Note The iSCSI hosts has not connected so they do not have a FCID 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)

  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)
```

Step 13 Check all iSCSI hosts to verify their connectivity.

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
  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 FC 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
  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
```


Send documentation comments to mdsfeedback-doc@cisco.com

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is ign.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: ign.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
```

Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Initiator ID based on IP address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FCIDs 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
ign.1987-05.com.cisco:01.25589167f74c]<-----
```

FCID resolved for iSCSI host

FCID for iSCSI host

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
```

Send documentation comments to mdsfeedback-doc@cisco.com

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      :
```

IP address of the iSCSI host

iSCSI gateway node

```
symbolic-node-name
```

```
:iqn.1987-05.com.cisco:01.25589167f74c<-----
```

iSCSI initiator ID is based on the registered node name

```
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      :
```

iSCSI gateway node

```
symbolic-node-name      :10.11.1.10  <-----
```

iSCSI initiator ID is based on the IP address registered in symbolic-node-name field

```
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :21:81:00:0b:fd:44:68:c0
hard-addr               :0x000000
```


[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Scenario 2

Sample scenario 2 (see [Figure 17-31](#)) assumes the following configuration:

- Access control based on Fibre Channel zoning.
- Target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator assigned to different VSANs.

Figure 17-31 iSCSI Scenario 2



To configure scenario 1 (see [Figure 17-31](#)), 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 FC 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 IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iscsi initiators by the IP address, and enable the interface.
- ```
switch(config)# int 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 IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

- 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)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

- Step 7** Add static configuration for each iSCSI initiator.

```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
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
switch(config-(iscsi-init))# static pwwn system-assigned 1
switch(config-(iscsi-init))# vsan 2
```

- Step 8** View the configured initiators.



**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.

- Based on the symbolic node name.

```
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
```

- 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
switch(config-zone)# member pwwn 21:00:00:20:37:6f:fd:97
```

- Step 11** Create 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
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 12** Activate the zoneset 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
```

**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<---- To FC target
 VSAN 1, ISID 00023d000001, Status active, no reservation
```

**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)<----- The configured nWWN
 Member of vsans: 1
 Number of Virtual n_ports: 1

 Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<----- The configured pWWN
 Interface iSCSI 7/1, Portal group tag: 0x300
 VSAN ID 1, FCID 0x680102
```

**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 isc..w <----- iSCSI initiator in
 name server
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 17** Verify the details of the iSCSI initiator's FCID 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

Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----- The configured nWWN
Member of vsans: 1
Number of Virtual n_ports: 1

Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<----- The configured pWWN
 Interface iSCSI 7/1, Portal group tag: 0x300
 VSAN ID 1, FCID 0x680102
```

**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 FCID 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
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 20** Verify that zoning has resolved the FCID 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** Do the same to 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 ign.1987-05.com.cisco:01.25589167f74c
Session #1
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <-- Session to
VSAN 2, ISID 00023d000001, Status active, no reservation first target

Session #2
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d <-- Session to
VSAN 2, ISID 00023d000001, Status active, no reservation second
target

switch# show iscsi initiator
iSCSI Node name is 10.15.1.11 <--- Initiator ID is the IP address
iSCSI Initiator name: ign.1987-05.com.cisco:01.25589167f74c
iSCSI alias name: oasis11.cisco.com

Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <----- Dynamic
Member of vsans: 2 <--- vsan membership WWN as
Number of Virtual n_ports: 1 static WWN
not
assigned

Virtual Port WWN is 20:06:00:0b:fd:44:68:c2 (configured) <----- Static
Interface iSCSI 7/5, Portal group tag: 0x304 pWWN for
VSAN ID 2, FCID 0x750200 the initiator

switch# show fcns database vsan 2
VSAN 2:

FCID TYPE PWWN (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 <-- iSCSI
Total number of entries = 3 initiator
name server
```

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

```
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] <-----
```

**FCID  
resolved for  
iSCSI  
initiator**

## Default IP Storage Settings

Table 17-2 lists the default settings for Gigabit Ethernet parameters.

**Table 17-2 Default Gigabit Ethernet Parameters**

| Parameters        | Default                           |
|-------------------|-----------------------------------|
| IP MTU frame size | 1500 bytes for all Ethernet ports |

Table 17-3 lists the default settings for FCIP parameters.

**Table 17-3 Default FCIP Parameters**

| Parameters                | Default            |
|---------------------------|--------------------|
| TCP default port for FCIP | 3225               |
| minimum-retransmit-time   | 300 milliseconds.  |
| keepalive-timeout         | 60 seconds.        |
| max-retransmissions       | 4 retransmissions. |
| PMTU discovery            | Enabled.           |
| pmtu-enable reset-timeout | 3600 seconds.      |
| SACK                      | Enabled.           |



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 17-3 Default FCIP Parameters**

| Parameters                              | Default                     |
|-----------------------------------------|-----------------------------|
| max-bandwidth                           | 1G.                         |
| min-available-bandwidth                 | 2 Mbps.                     |
| round-trip-time                         | 10ms.                       |
| buffer size                             | 0 KB.                       |
| Control TCP and data connection         | No packets are transmitted. |
| TCP congestion window monitoring        | Enabled                     |
| Burst size                              | 10KB.                       |
| TCP connection mode                     | active mode is enabled.     |
| special-frame                           | Disabled.                   |
| FCIP timestamp                          | Disabled.                   |
| acceptable-diff range to accept packets | + or - 1000 milliseconds.   |
| B port keepalive responses              | Disabled                    |

Table 17-4 lists the default settings for iSCSI parameters.

**Table 17-4 Default iSCSI Parameters**

| Parameters                                         | Default                                                                                                             |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Number of TCP connections                          | One per iSCSI session.                                                                                              |
| 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 default VSAN (VSAN 1).                                                                               |
| Identifying initiators                             | iSCSI node names.                                                                                                   |
| Advertising static virtual targets                 | No initiators allowed to access a virtual target (unless explicitly configured).                                    |
| iSCSI login authentication                         | CHAP or none authentication mechanism.                                                                              |
| Ethernet PortChannel IP address usage              | Source and destination IP addresses.                                                                                |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***



# CHAPTER 18

## 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.

This chapter provides configuration and messaging details on the Call Home feature. It includes the following sections:

- [Call Home Features, page 18-2](#)
- [Call Home Configuration Process, page 18-2](#)
- [Cisco AutoNotify, page 18-3](#)
- [Configuring the Call Home Function, page 18-3](#)
- [Assigning Contact Information, page 18-4](#)
- [Configuring Destination Profiles, page 18-5](#)
- [Configuring E-Mail Options, page 18-6](#)
- [Enabling or Disabling Call Home, page 18-7](#)
- [Testing Call Home Communication, page 18-8](#)
- [Displaying Call Home Information, page 18-8](#)
- [Default Settings, page 18-9](#)
- [Event Triggers, page 18-10](#)
- [Call Home Message Severity Levels, page 18-11](#)
- [Message Contents, page 18-12](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family. It provides multiple Call Home profiles, each with separate potential destinations.

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 Connection Online (CCO) website at <http://www.cisco.com/>. The XML format enables communication with the Cisco Systems TAC group.
- Multiple concurrent message destinations. Up to 50 E-mail destination addresses are allowed for each format type.
- Message categories include system, environment, switching module hardware, supervisor module, hardware, inventory, and test.

## Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- E-mail server and at least one destination profile must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, email, or automated service such as Cisco AutoNotify.
- 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 for the feature to operate.
- If Cisco AutoNotify is used, an active service contract must cover the device being configured.

To configure Call Home, follow these steps:

- 
- |               |                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Configure the Call Home function (see the “ <a href="#">Configuring the Call Home Function</a> ” section on page 18-3). |
| <b>Step 2</b> | Assign contact information (see the “ <a href="#">Assigning Contact Information</a> ” section on page 18-4).            |
| <b>Step 3</b> | Configure destination profiles (see the “ <a href="#">Configuring Destination Profiles</a> ” section on page 18-5).     |
| <b>Step 4</b> | Enable or disable Call Home (see the “ <a href="#">Enabling or Disabling Call Home</a> ” section on page 18-7).         |
| <b>Step 5</b> | Test Call Home messages (see the “ <a href="#">Testing Call Home Communication</a> ” section on page 18-8).             |
-

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible through registration with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support. To register, the following items are required:

- The SMARTnet contract number covering your MDS 9000 family switch.
- Your name, company address, your email address, and your CCO ID.
- 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), or by executing the operating system **show sprom backplane 1** command.
- The exact product number of your Cisco MDS 9000 Family switch. This can be obtained by executing the same operating system command as above. For example, some valid product numbers include: DS-C6509 and DS-C9216-K9

To configure a Cisco MDS 9000 Family switch to use AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and email address information is found on the Cisco.com web site at:

[http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/annoti\\_ds.htm](http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/annoti_ds.htm)



### Note

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, email server, and an XML destination profile as specified in the Service Activation document ([http://www.cisco.com/univercd/cc/td/doc/product/voice/c\\_callmg/3\\_3/service/serv332/ccm/srvs/sssrvtact.htm](http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccm/srvs/sssrvtact.htm)). The **contract-id**, **customer-id**, **site-id**, and **switch-priority** parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

## Configuring the Call Home Function

To enter the Call Home submode, follow these steps:

|        | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Purpose                                        |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Enters configuration mode.                     |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Enters Call Home submode.                      |
| Step 3 | switch(config-callhome)# <b>?</b><br>contract-id           Service contract id of the customer<br>customer-id           Customer id<br>destination-profile   Configure destination profiles<br>disable               Disable callhome<br>email-contact        Email address of the contact person<br>enable                Enable callhome<br>exit                  Exit from this submode<br>no                    Negate a command or set its defaults<br>phone-contact        Contact person's phone number<br>site-id                Site id of the network where switch is deployed<br>streetaddress        Configure replacement part shipping address.<br>switch-priority       Priority of the switch(0-highest 7-lowest)<br>transport            Configure transport related configuration | Displays the options available at this prompt. |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Assigning Contact Information

It is mandatory for each switch to include e-mail, phone, and street address information. It's 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# <b>config t</b>                                                                                                                                                     | Enters configuration mode.                                                                                                                                                                          |
| Step 2  | switch# <b>snmp-server contact</b><br>personname@companyname.com                                                                                                            | Configures the SNMP contact e-mail address to receive a test message reply from Cisco.                                                                                                              |
| Step 3  | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                                                 | Enters the Call Home submode.                                                                                                                                                                       |
| Step 4  | switch(config-callhome)# <b>email-contact</b><br><b>username@company.com</b><br>successfully updated the information<br>switch(config-callhome)#                            | Assigns the customer's e-mail address. Up to 128 alphanumeric characters are accepted in e-mail address format.<br><br><b>Note</b> You can use any valid e-mail address. You cannot use spaces.     |
| Step 5  | switch(config-callhome)# <b>phone-contact</b><br><b>+1-800-123-4567</b><br>successfully updated the information<br>switch(config-callhome)#                                 | Assigns the customer's phone number. Up to 20 alphanumeric characters are accepted in international format.<br><br><b>Note</b> You cannot use spaces. Be sure to use the + prefix before the number |
| Step 6  | switch(config-callhome)# <b>streetaddress 1234</b><br><b>Picaboo Street, Any city, Any state, 12345</b><br>successfully updated the information<br>switch(config-callhome)# | 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)# <b>switch-priority 0</b><br>successfully updated the information<br>switch(config-callhome)#                                                       | Assigns the switch priority, with 0 being the highest priority and 7 the lowest.<br><br><b>Tip</b> Use this field to create a hierarchical management structure.                                    |
| Step 8  | switch(config-callhome)# <b>customer-id</b><br><b>Customer1234</b><br>successfully updated the information<br>switch(config-callhome)#                                      | Optional. Identifies the customer ID. Up to 256 alphanumeric characters are accepted in free format.                                                                                                |
| Step 9  | switch(config-callhome)# <b>site-id</b><br><b>Site1ManhattanNY</b><br>successfully updated the information<br>switch(config-callhome)#                                      | Optional. Identifies the customer site ID. Up to 256 alphanumeric characters are accepted in free format.                                                                                           |
| Step 10 | switch(config-callhome)# <b>contract-id</b><br><b>Company1234</b><br>successfully updated the information<br>switch(config-callhome)#                                       | Assigns the customer ID for the switch. Up to 64 alphanumeric characters are accepted in free format.                                                                                               |

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Configuring 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.



### Note

If you use the Cisco AutoNotify service, the XML destination profile is required (see [http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti\\_ds.htm](http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti_ds.htm)).

- **Profile ID**—A text string that uniquely identifies three predefined destination profile formats: full text, short text, and XML.
- **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 destination profile messaging options, follow these steps:

|        | Command                                                                                                                                                                                          | Purpose                                                                                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                          | Enters configuration mode.                                                                                                                                                                                                                       |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                                                                      | Enters the Call Home submode.                                                                                                                                                                                                                    |
| Step 3 | switch(config-callhome)#<br><b>destination-profile</b><br><b>full-txt-destination email-addr</b><br><b>person@place.com</b><br>successfully updated the information<br>switch(config-callhome)#  | Configures a destination e-mail address for a message sent in full text format. This text provides the complete, detailed explanation of the failure.<br><b>Tip</b> Use a standard e-mail address that does not have any text size restrictions. |
|        | switch(config-callhome)#<br><b>destination-profile</b><br><b>full-txt-destination message-size</b><br><b>1000000</b><br>successfully updated the information<br>switch(config-callhome)#         | Configures a destination message size for a message sent in full text 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 4 | switch(config-callhome)#<br><b>destination-profile</b><br><b>short-txt-destination email-addr</b><br><b>person@place.com</b><br>successfully updated the information<br>switch(config-callhome)# | Configures a destination e-mail address for a message sent in short text format. This text provides the basic explanation of the failure.<br><b>Tip</b> Use a pager-related e-mail address for this option.                                      |
|        | switch(config-callhome)#<br><b>destination-profile</b><br><b>short-txt-destination message-size</b><br><b>100000</b><br>successfully updated the information                                     | Configures a destination message size for a message sent in short text format. 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.                                     |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

|        | Command                                                                                                                                                            | Purpose                                                                                                                                                                                                                                                                                           |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 5 | <pre>switch(config-callhome)# destination-profile XML-destination email-addr findout@cisco.com successfully updated the information switch(config-callhome)#</pre> | Configures a destination e-mail address for a message sent in XML format. This option provides the full information that is compatible with Cisco Systems TAC support.<br><br><b>Tip</b> Do not add a pager-related e-mail address to this destination profile because of the large message size. |
|        | <pre>switch(config-callhome)# destination-profile XML-destination message-size 100000 successfully updated the information</pre>                                   | Configures a destination message size for a message sent in 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.                                                                                          |

**Note**

Steps 3, 4, and 5 in this procedure can be skipped or configured in any order.

## Configuring 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 ensure to configure the SMTP server address and port number for the Call Home functionality to work.

## Configuring General E-Mail Option

To configure general e-mail options, follow these steps:

|        | Command                                                                                                                                     | Purpose                                                                                 |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Step 1 | <pre>switch# config t</pre>                                                                                                                 | Enters configuration mode.                                                              |
| Step 2 | <pre>switch(config)# callhome switch(config-callhome)#</pre>                                                                                | Enters Call Home submode.                                                               |
| Step 3 | <pre>switch(config-callhome)# transport email from user@company1.com successfully updated the information switch(config-callhome)#</pre>    | Optional. Configures the from e-mail address.                                           |
| Step 4 | <pre>switch(config-callhome)# transport email reply-to person@place.com successfully updated the information switch(config-callhome)#</pre> | Optional. Configures the reply-to e-mail address to which all responses should be sent. |



*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Configuring SMTP Server and Ports

To configure the SMTP server and port, follow these steps:

|        | Command                                                                                                                                             | Purpose                                                                                                                         |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                             | Enters configuration mode.                                                                                                      |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                         | Enters Call Home submode.                                                                                                       |
| Step 3 | switch(config-callhome)# <b>transport email smtp-server 192.168.1.1</b><br>successfully updated the information<br>switch(config-callhome)#         | Configures the DNS or IP address of the SMTP server to reach the server. The port usage defaults to 25 if no port is specified. |
|        | switch(config-callhome)# <b>transport email smtp-server 192.168.1.1 port 30</b><br>successfully updated the information<br>switch(config-callhome)# | <b>Note</b> The port number is optional and, if required, may be changed depending on the server location.                      |

## Enabling or Disabling Call Home

Once you have configured the contact information, you must enable the Call Home function. The **enable** command is required for the Call Home function to start operating.

To enable the Call Home function, follow these steps:

|        | Command                                                                                             | Purpose                                                                                                  |
|--------|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                             | Enters configuration mode.                                                                               |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                         | Enters Call Home submode.                                                                                |
| Step 3 | switch(config-callhome)# <b>enable</b><br>callhome enabled successfully<br>switch(config-callhome)# | Enables the Call Home function.                                                                          |
|        | switch(config-callhome)# <b>disable</b><br>switch(config-callhome)#                                 | Disables the Call Home function. When you disable the Call Home function, all input events are ignored.  |
|        |                                                                                                     | <b>Note</b> Even if Call Home is disabled, basic information for each Call Home event is sent to syslog. |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Testing Call Home Communication

You can simulate a message generation by issuing a **test** command.

To test the Call Home function, follow these steps:

|        | Command                                                                                                                              | Purpose                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Step 1 | switch# <b>callhome test</b><br>trying to send test callhome message<br>successfully sent test callhome message<br>switch#           | Sends a test message to the configured destination(s).           |
| Step 2 | switch# <b>callhome test inventory</b><br>trying to send test callhome message<br>successfully sent test callhome message<br>switch# | 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 18-1 to 18-6).

### Example 18-1 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:Andiamo1234
switch priority:0
```

### Example 18-2 Displays Destination Profile Information

```
switch# show callhome destination-profile
XML destination profile information
maximum message size:250000
email addresses configured:
findout@cisco.com

Short-txt destination profile information
maximum message size:4000
email addresses configured:
person1@epage.company.com

full-txt destination profile information
maximum message size:250000
email addresses configured:
person2@company2.com
```

### Example 18-3 Displays the Full-Text Profile

```
switch# show callhome destination-profile profile full-txt-destination
full-txt destination profile information
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```
maximum message size:250000
email addresses configured:
person2@company2.com
```

**Example 18-4 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 18-5 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 18-6 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
```

## Default Settings

Table 18-1 lists the default Call Home default settings.

**Table 18-1 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.                 | 4,000   |
| DNS or IP address of the SMTP server to reach the server if no port is specified. | 25      |

**[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned commands to execute when the event occurs. The command output is included in the transmitted message. [Table 18-2](#) lists the trigger events. [Table 18-3](#) lists event categories and command outputs.

**Table 18-2 Event Triggers**

| Event     | Type              | Event Name                   | Description                                                                                 | Severity Level |
|-----------|-------------------|------------------------------|---------------------------------------------------------------------------------------------|----------------|
| Call Home | System            | SW_CRASH_STATELESS_RESTART   | Software process crashed with a disruptive restart indicating an interruption of a service. | 5              |
|           |                   |                              |                                                                                             |                |
|           | Environmental     | 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  | LINECARD_FAILURE             | Switching module operation failed.                                                          | 7              |
|           |                   | POWER_UP_DIAGNOSTICS_FAILURE | Switching module failed power up diagnostics.                                               | 7              |
|           | Supervisor module | SUP_FAILURE                  | Supervisor module operation failed.                                                         | 7              |
|           |                   | POWER_UP_DIAGNOSTICS_FAILURE | Supervisor module failed power up diagnostics.                                              | 7              |
| Inventory | Inventory         | 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              | TEST                         | User generated test.                                                                        | 2              |

**Table 18-3 Event Categories and Command Outputs**

| Event Category | Description                                                                                | Executed Commands                                                |
|----------------|--------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| System         | Events generated by failure of a software system that is critical to unit operation.       | <b>show tech-support</b><br><b>show system redundancy status</b> |
| Environmental  | Events related to power, fan, and environment sensing elements such as temperature alarms. | <b>show module</b><br><b>show environment</b>                    |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Table 18-3 Event Categories and Command Outputs (continued)**

| Event Category            | Description                                                                                                                                                                                        | Executed Commands        |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| Switching module hardware | Events related to standard or intelligent switching modules.                                                                                                                                       | <b>show tech-support</b> |
| Supervisor hardware       | Events related to supervisor modules.                                                                                                                                                              | <b>show tech-support</b> |
| 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. | <b>show version</b>      |
| Generic                   | User definable Call Home triggers.                                                                                                                                                                 | <b>show tech-support</b> |
| Test                      | User generated test message.                                                                                                                                                                       | <b>show version</b>      |

## Call Home Message Severity Levels

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Severity levels are preassigned per event type.



### Note

Call Home severity levels are not the same as system message logging severity levels (see [Chapter 21, “Configuring System Message Logging”](#)).

Severity levels range from 0 to 9, with 9 having the highest urgency. Each severity level has keywords as listed in [Table 18-4](#).

**Table 18-4 Severity Levels**

| Severity Level | Keyword      | Description                                                                          |
|----------------|--------------|--------------------------------------------------------------------------------------|
| 9              | Catastrophic | Network wide catastrophic failure.                                                   |
| 8              | Disaster     | Significant network impact.                                                          |
| 7              | Fatal        | System is unusable.                                                                  |
| 6              | Critical     | Critical conditions, immediate attention needed.                                     |
| 5              | Major        | Major conditions.                                                                    |
| 4              | Minor        | Minor conditions.                                                                    |
| 3              | Warning      | Warning conditions.                                                                  |
| 2              | Notification | Basic notification and informational messages. Possibly independently insignificant. |
| 1              | Normal       | Normal event signifying return to normal state.                                      |
| 0              | Debugging    | Debugging messages.                                                                  |

**[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## 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 18-5](#) describes the short text formatting option for all message types.

**Table 18-5 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 syslog message |

[Table 18-6](#), [Table 18-7](#), and [Table 18-8](#) display the information contained in plain text and XML messages.

**Table 18-6 Reactive Event Message Format**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                               | XML Tag<br>(XML only) |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Time stamp                        | Date and time stamp of event in ISO time notation:<br><i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> 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 <a href="#">“Event Triggers”</a> section on page 18-10.                                                                                                                                   | /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 <a href="#">Table 18-4</a> ).                                                                                                                                                                                      | /mml/header/level     |
| Source ID                         | Product type for routing.                                                                                                                                                                                                                         | /mml/header/source    |

## Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

**Table 18-6 Reactive Event Message Format (continued)**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | XML Tag<br>(XML only)            |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Device ID                         | <p>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</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /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                         | <p>If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /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        |
| 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      |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 18-6 Reactive Event Message Format (continued)**

| Data Item<br>(Plain text and XML)  | Description<br>(Plain text and XML)                                          | XML Tag<br>(XML only)              |
|------------------------------------|------------------------------------------------------------------------------|------------------------------------|
| 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                | Exact command that was run. For example, <b>show running-config</b> 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 <a href="#">Table 18-3</a> ).  | /mml/attachments/attachment/atdata |

**Table 18-7 Inventory Event Message Format**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                               | XML Tag<br>(XML only) |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Time stamp                        | Date and time stamp of event in ISO time notation:<br><i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> 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 “ <a href="#">Event Triggers</a> ” section on page 18-10.                                                                                                 | /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 (see <a href="#">Table 18-4</a> ).                                                                                                                                                                   | /mml/header/level     |
| Source ID                         | Product type for routing at Cisco. Specifically “MDS 9000”                                                                                                                                                                                        | /mml/header/source    |



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 18-7 Inventory Event Message Format (continued)**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>XML Tag<br/>(XML only)</b>    |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Device ID                                 | <p>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</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /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                                 | <p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /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        |
| 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      |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 18-7 Inventory Event Message Format (continued)**

| Data Item<br>(Plain text and XML)  | Description<br>(Plain text and XML)                                                                                            | XML Tag<br>(XML only)              |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| 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                | Exact command that was run. For example, the <b>show running-config</b> 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 <a href="#">“Event Triggers” section on page 18-10</a> ). | /mml/attachments/attachment/atdata |

**Table 18-8 User-Generated Test Message Format**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                            | XML Tag<br>(XML only) |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Time stamp                        | Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> 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 <a href="#">“Event Triggers” section on page 18-10</a> .                                                                                  | /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 <a href="#">Table 18-4</a> ).                                                                                                                                                           | /mml/header/level     |
| Source ID                         | Product type for routing.                                                                                                                                                                                                                      | /mml/header/source    |

## Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

**Table 18-8 User-Generated Test Message Format (continued)**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | XML Tag<br>(XML only)                  |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| Device ID                         | <p>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</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /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                         | <p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /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                   |
| 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 <a href="#">Table 18-3</a> .                                                                                                                                                                                                                                                                                                                                                                                                           | /mml/attachments/attachmen<br>t/atdata |
| MIME type                         | Normally text or plain or encoding type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | /mml/attachments/attachmen<br>t/mime   |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 18-8** *User-Generated Test Message Format (continued)*

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                      | <b>XML Tag<br/>(XML only)</b>        |
|-------------------------------------------|----------------------------------------------------------------------------------|--------------------------------------|
| Attachment type                           | Specifically command output.                                                     | /mml/attachments/attachmen<br>t/type |
| Command output<br>name                    | Exact command that was run. For example, the <b>show running-config</b> command. | /mml/attachments/attachmen<br>t/name |



## 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.

This chapter includes the following sections:

- [About fcdomain Phases, page 19-2](#)
- [Restarting the Domain, page 19-3](#)
- [Configuring the Domain, page 19-4](#)
- [Setting Switch Priority, page 19-6](#)
- [Merging Stable Fabrics, page 19-6](#)
- [Assigning Contiguous Domains, page 19-7](#)
- [Disabling the fcdomain Feature, page 19-7](#)
- [Setting the Fabric Name, page 19-8](#)
- [Stopping Incoming RCFs, page 19-8](#)
- [Enabling Persistent FC IDs, page 19-9](#)
- [Configuring Persistent FC IDs Manually, page 19-10](#)
- [Purging Persistent FC IDs, page 19-11](#)
- [Displaying fcdomain Information, page 19-12](#)
- [Default Settings, page 19-14](#)



### 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 using the **copy running-config startup-config** command. 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.

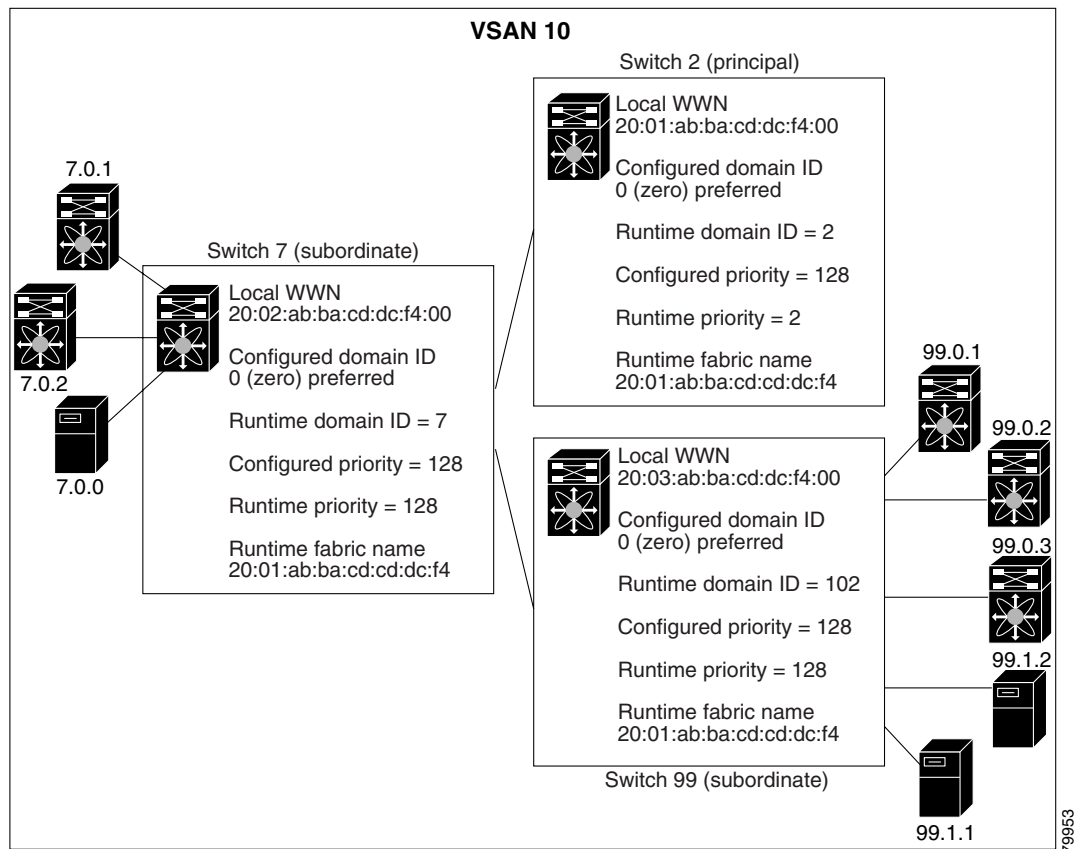
**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## About fcdomain Phases

This section describes each fcdomain phase (see [Figure 19-1](#)):

- 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.

**Figure 19-1 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.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Restarting the Domain

The **fcdomain restart** command applies your changes to the runtime settings. 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. If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric.

To restart the fabric disruptively or nondisruptively, follow these steps:

|        | Command                                                   | Purpose                                                      |
|--------|-----------------------------------------------------------|--------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                | Enters configuration mode.                                   |
| Step 2 | switch(config)# <b>fcdomain restart vsan 1</b>            | Forces the VSAN to reconfigure without traffic disruption.   |
|        | switch(config)# <b>fcdomain restart disruptive vsan 1</b> | Forces the VSAN to reconfigure with data traffic disruption. |

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



### 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 will change to take on the static domain ID after the next restart (see the [“Configuring the Domain”](#) section on page 19-4).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

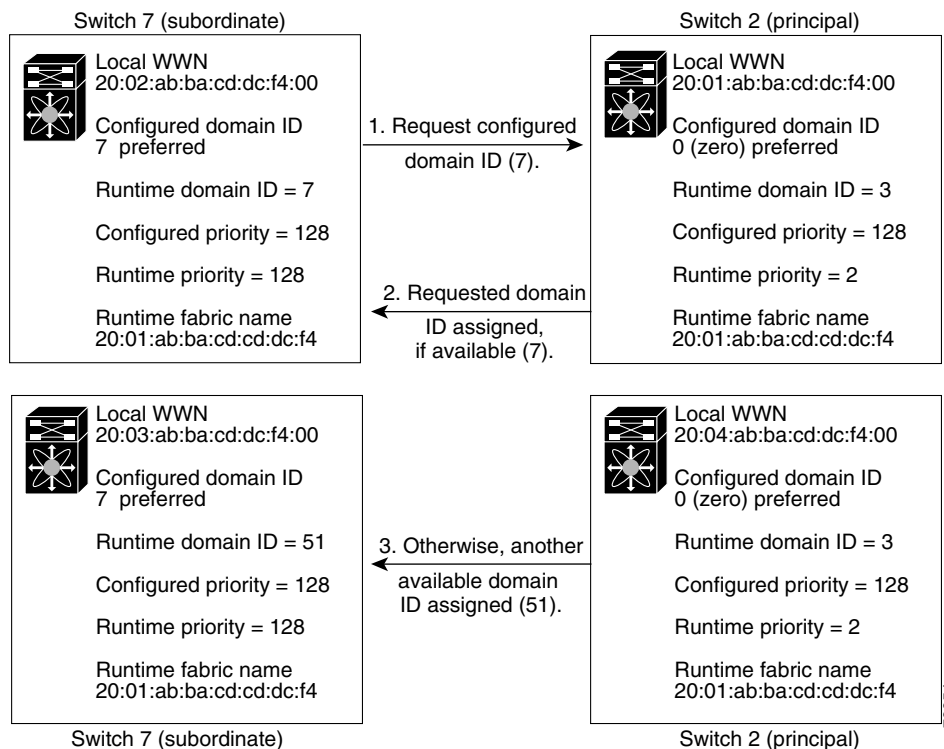
## Configuring the Domain

The configured domain ID can be **preferred** or **static**. By default, the configured domain is **0** and the configured option is **preferred**. If you do not configure a domain ID, the local switch sends a random ID in its request.

When a subordinate switch requests a domain, the following process takes place (see Figure 19-2):

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.
3. Otherwise, it assigns another available domain ID.

**Figure 19-2 Configuration Process Using the preferred Option**



A subordinate switch behavior changes based on the option of its configured domain ID and the domain ID that the principal switch has assigned to the requesting switch:

- 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 option 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 option is **preferred**, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Caution**

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

To specify a **preferred** or a **static** domain ID, follow these steps:

|               | Command                                                      | Purpose                                                                                                                                                                      |
|---------------|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b><br>switch(config)#                   | Enters configuration mode.                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>fcdomain domain 3 preferred vsan 8</b>    | Configures the switch in VSAN 8 to request a preferred domain ID 3 and accepts any value assigned by the principal switch.                                                   |
|               | switch(config)# <b>no fcdomain domain 3 preferred vsan 8</b> | Resets the configured domain ID to 0 (default) in VSAN 8. The configured domain ID becomes 0 preferred.                                                                      |
| <b>Step 3</b> | switch(config)# <b>fcdomain domain 2 static vsan 237</b>     | 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)# <b>no fcdomain domain 18 static vsan 237</b> | Resets the configured domain ID to factory defaults in VSAN 237. The configured domain ID becomes 0 preferred.                                                               |

**Note**

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

While the **static** option can be applied to runtime after a disruptive or nondisruptive restart, the **preferred** option is applied to runtime only after a disruptive restart (see the [“Restarting the Domain” section on page 19-3](#)).

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Setting 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.

To configure the priority for the principal switch, follow these steps:

|        | Command                                                | Purpose                                                       |
|--------|--------------------------------------------------------|---------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#             | Enters configuration mode.                                    |
| Step 2 | switch(config)# <b>fcdomain priority 25 VSAN 99</b>    | Configures a priority of 25 for the local switch in VSAN 99.  |
|        | switch(config)# <b>no fcdomain priority 25 VSAN 99</b> | Reverts the priority to the factory default (128) in VSAN 99. |

The priority configuration is applied to runtime through a disruptive restart (see the [“Restarting the Domain”](#) section on page 19-3).

## Merging Stable Fabrics

By default, the **auto-reconfigure** option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

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

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

|        | Command                                                  | Purpose                                                                                         |
|--------|----------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#               | Enters configuration mode.                                                                      |
| Step 2 | switch(config)# <b>fcdomain auto-reconfigure vsan 10</b> | Enables the automatic reconfiguration option in VSAN 10.                                        |
|        | switch(config)# <b>no fcdomain auto-reconfigure 69</b>   | Disables the automatic reconfiguration option and reverts it to the factory default in VSAN 69. |

The **auto-reconfigure** 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 **auto-reconfigure** option on both switches, the fabric continues to be isolated. However, if you enable the **auto-reconfigure** option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. A disruptive reconfiguration may affect data traffic. You can nondisruptively perform this function by changing the configured domains on the overlapping links and getting rid of the overlaps.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Assigning Contiguous Domains

By default, the **contiguous-allocation** option is disabled. When the subordinate switches request the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the **contiguous-allocation** option is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches.
- If the **contiguous-allocation** option is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switches.

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

|        | Command                                                            | Purpose                                                                                       |
|--------|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                         | Enters configuration mode.                                                                    |
| Step 2 | switch(config)# <b>fcdomain contiguous-allocation vsan 81-83</b>   | Enables the contiguous allocation option in VSAN 81 through 83.                               |
|        | switch(config)# <b>no fcdomain contiguous-allocation vsan 1030</b> | Disables the contiguous allocation option and reverts it to the factory default in VSAN 1030. |

The **contiguous-allocation** option takes immediate effect at runtime—you do not need to restart the fcdomain.

## Disabling the fcdomain Feature

By default, the fcdomain feature is enabled on each switch. You can disable the fcdomain feature by using the **no fcdomain** command. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric.

To disable fcdomains in a single VSAN or a range of VSANs, follow these steps:

|        | Command                                       | Purpose                                                    |
|--------|-----------------------------------------------|------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#    | Enters configuration mode.                                 |
| Step 2 | switch(config)# <b>no fcdomain vsan 7-200</b> | Disables the fcdomain configuration in VSAN 7 through 200. |
|        | switch(config)# <b>fcdomain vsan 2008</b>     | Enables the fcdomain configuration in VSAN 2008.           |

The fcdomain configuration is applied to runtime through a disruptive restart.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Setting the Fabric Name

By default the configured fabric name is 20:01:00:05:30:00:28:df.

- When the `fcdomain` feature is disabled, the runtime fabric name is the same as the configured fabric name.
- When the `fcdomain` feature is enabled, the runtime fabric name is the same as the principal switch's WWN.

To set the fabric name value for a disabled `fcdomain`, follow these steps:

|        | Command                                                                                  | Purpose                                                                                      |
|--------|------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                               | Enters configuration mode.                                                                   |
| Step 2 | switch(config)# <b>fcdomain fabric-name</b><br><b>20:1:ac:16:5e:0:21:01 vsan 3</b>       | Assigns the configured fabric name value in VSAN 3.                                          |
|        | switch(config)# <b>no fcdomain fabric-name</b><br><b>20:1:ac:16:5e:0:21:01 vsan 3010</b> | Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010. |

The fabric name is applied to runtime through a disruptive restart when the `fcdomain` is configured as disabled (see the [“Restarting the Domain”](#) section on page 19-3).

## Stopping Incoming RCFs

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

To stop incoming RCF request frames, follow these steps:

|        | Command                                                 | Purpose                                                       |
|--------|---------------------------------------------------------|---------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#              | Enters configuration mode.                                    |
| Step 2 | switch(config)# <b>int fc1/1</b><br>switch(config-if)#  | Configures the specified interface.                           |
| Step 3 | switch(config-if)# <b>fcdomain rcf-reject vsan 1</b>    | Enables the RCF filter on the specified interface in VSAN 1.  |
|        | switch(config-if)# <b>no fcdomain rcf-reject vsan 1</b> | Disables the RCF filter on the specified interface in VSAN 1. |

The **rcf-reject** option takes immediate effect to runtime through a disruptive restart (see the [“Restarting the Domain”](#) section on page 19-3).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Enabling Persistent FC IDs

When a N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned a FC ID. By default, the persistent FC ID feature is disabled. If this feature is disabled, the following consequences apply:

- A 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 it was originally connected.

The assigned FC IDs in a fcdomain can be enabled to remain persistent even after a reboot. This ensures that an attached N port receives the same FC IDs after a reboot. If you enable this feature, 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.

To enable the persistent FC ID feature, you must meet the following requirements:

- Configure a static domain ID in that VSAN.
- Ascertain that the static configured domain is same as the runtime domain. You can verify if they are the same by issuing the show fcdomain command



**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.

To enable the persistent FC ID feature, follow these steps:

|        | Command                                                                                             | Purpose                                                             |
|--------|-----------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                             | Enters configuration mode.                                          |
| Step 2 | switch(config)# <b>fcdomain domain 2 static vsan 1000</b>                                           | Configures the switch in VSAN 1000 to accept only a specific value. |
| Step 3 | switch(config)# <b>fcdomain fcid persistent vsan 1000</b><br>FCID(s) persistent feature is enabled. | Activates persistency of FC IDs in VSAN 1000.                       |
|        | switch(config-if)# <b>no fcdomain fcid persistent vsan 20</b>                                       | Disables the FC ID persistency feature in VSAN 20.                  |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Note**

Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID. You can enable the FC ID non-disruptively by configuring the static domain and the runtime domain to be the same. You can obtain the runtime domain by issuing the **show fcdomain** command.

To enable the persistent FC ID feature, follow these steps:

|        | Command                                                                                                                       | Purpose                                                                                                                                                                                                                                                                 |
|--------|-------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                                                    | Enters configuration mode.                                                                                                                                                                                                                                              |
| Step 2 | switch(config)# <b>fcdomain domain 2 static vsan 55</b>                                                                       | Configures the switch in VSAN 55 to accept only a specific value (2 in this example).<br><br><b>Note</b> Once the VSAN is restarted, if the requested domain ID is not granted by the principal switch, all local interfaces in VSAN 55 are moved to an isolated state. |
| Step 3 | switch(config)# <b>fcdomain fcid persistent vsan 55</b><br>ERROR: Static Configured Domain ID and Runtime Domain ID mismatch. | Returns an error due to a domain ID mismatch in VSAN 55.                                                                                                                                                                                                                |
| Step 4 | switch(config)# <b>fcdomain restart vsan 55</b>                                                                               | Forces the VSAN to reconfigure without traffic disruption.                                                                                                                                                                                                              |
| Step 5 | switch(config)# <b>fcdomain fcid persistent vsan 55</b><br>FCID(s) persistent feature is enabled.                             | Activates persistency of FC IDs in VSAN 55.                                                                                                                                                                                                                             |

## Configuring Persistent FC IDs Manually

Once 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 require 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.
- Do not replace an FC ID that is already configured in another WWN. If you want to use a previously-configured WWN, first delete the configured WWN before proceeding with this procedure.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

To configure persistent FC IDs, follow these steps:

|        | Command                                                                                    | Purpose                                                                                                     |
|--------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                 | Enters configuration mode.                                                                                  |
| Step 2 | switch(config)# <b>fcdomain fcid database</b>                                              | Activates persistency of FC IDs in the specified VSAN.                                                      |
| Step 3 | switch(config-fcid-db)# <b>vsan 1000 wwn 33:e8:00:05:30:00:16:df fcid 0x070128</b>         | Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in VSAN 1000.                     |
|        | switch(config-fcid-db)# <b>vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070123 dynamic</b> | 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)# <b>vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070100 area</b>    | Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC IDs 0x070100 through 0x0701FF in VSAN 1000.   |
|        |                                                                                            | <b>Note</b> To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID. |

## Purging Persistent FC IDs

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. [Table 19-1](#) identifies the FC ID entries that are deleted by the **purge fcdomain** command.

**Table 19-1 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     |

Dynamic, not in use, FC IDs can be removed using the **purge fcdomain** command (see [Table 19-1](#)).

To purge persistent FC IDs, follow this step:

|        | Command                                     | Purpose                                                   |
|--------|---------------------------------------------|-----------------------------------------------------------|
| Step 1 | switch# <b>purge fcdomain fcid vsan 4</b>   | Purges all dynamic and unused FC IDs in VSAN 4            |
|        | switch# <b>purge fcdomain fcid vsan 3-5</b> | Purges all dynamic and unused FC IDs in VSAN 3, 4, and 5. |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Displaying fcdomain Information

The **show fcdomain** commands display global information about the fcdomain configurations. See [Example 19-1](#).



### Note

In [Example 19-1](#), the fcdomain feature is disabled. Consequently, the runtime fabric name is the same as the configured fabric name.

### Example 19-1 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 **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. See [Example 19-2](#).

### Example 19-2 Displays the fcdomain List

```
switch# show fcdomain domain-list vsan 1

Number of domains: 1
Domain ID WWN

0x16(22) 20:01:00:05:30:00:16:df [Local] [Principal]
```

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 19-3](#) and [19-4](#).

### Example 19-3 Displays Persistent FC IDs in a Specified VSAN

```
switch# show fcdomain fcid persistent vsan 1000
Total entries 2.

Persistent FCIDs table contents:
```



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

| 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 19-4 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 19-5](#) and [Example 19-6](#).

#### Example 19-5 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 19-6 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 19-7](#).

#### Example 19-7 Displays FC ID Information

```
switch# show fcdomain address-allocation vsan 1
Free FCIDs: 0x650108 to 0x65ffff

Assigned FCIDs: 0x650000 to 0x650107

Reserved FCIDs: 0x65ffff
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```
Number free FCIDs: 65271
Number assigned FCIDs: 264
Number reserved FCIDs: 1
```

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 19-8](#).

**Example 19-8 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 19-2](#) lists the default settings for all fcdomain parameters.

**Table 19-2 Default fcdomain Parameters**

| Parameters                   | Default                  |
|------------------------------|--------------------------|
| fcdomain feature             | Enabled.                 |
| Configured domain ID         | 0 (zero).                |
| Configured domain option     | Preferred.               |
| auto-reconfigure option      | Disabled.                |
| contiguous-allocation option | Disabled.                |
| Priority                     | 128.                     |
| Fabric-name                  | 20:01:00:05:30:00:28:df. |
| rcf-reject                   | Disabled.                |
| Persistent FC ID             | Disabled.                |



## Configuring Traffic Management

---

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:

- Guarantees absolute and relative bandwidth choices
- Provides latency to reduce frame loss
- Prioritizes transactional traffic over bulk traffic
- Supports multiple VSANs on the same fabric by guaranteeing bandwidth and latency available to each VSAN

This chapter provides details on the QoS and FCC features provided in all switches. It includes the following sections:


- [FCC, page 20-2](#)
- [QoS, page 20-4](#)
- [Default FCC and QoS Settings, page 20-4](#)

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see [Figure 20-1](#)).

**Figure 20-1 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).

## FCC Process

When a node in the network detects a 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 quest 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 quest frames. However, only the edge switch processes edge quest frames. The FCC protocol is implemented for each VSAN and can be enabled or disabled on a specified VSAN or for all VSANs at the same time.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled globally for the entire switch.

To enable or disable the FCC feature, follow these steps:

|        | Command                       | Purpose                    |
|--------|-------------------------------|----------------------------|
| Step 1 | switch# <b>config t</b>       | Enters configuration mode. |
| Step 2 | switch(config)# <b>fcc</b>    | Enables FCC globally.      |
|        | switch(config)# <b>no fcc</b> | Disables FCC globally.     |

## Assigning FCC Priority

To assign FCC priority, follow these steps:

|        | Command                               | Purpose                                                                             |
|--------|---------------------------------------|-------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>               | Enters configuration mode.                                                          |
| Step 2 | switch(config)# <b>fcc priority 2</b> | Defines the FCC priority threshold with 0 being the lowest and 7 being the highest. |

## Displaying FCC

Use the show fcc commands to view FCC settings (see [Example 20-1](#)).

### ***Example 20-1 Displays Configured FCC Information***

```
switch# show fcc
fcc is disabled
fcc is applied to frames with priority up to 4
```

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

# QoS



**Note**

In Release 1, the QoS functionality provides control traffic over data 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 the control traffic is enabled (priority 0 is the default).

To disable the high priority assignment for control traffic, follow these steps:

|        | Command                                                             | Purpose                                   |
|--------|---------------------------------------------------------------------|-------------------------------------------|
| Step 1 | switch# <b>config t</b>                                             | Enters configuration mode.                |
| Step 2 | switch(config)# <b>no qos control priority 0</b><br>switch(config)# | Disables the control traffic QoS feature. |
| Step 3 | switch(config)# <b>qos control priority 0</b><br>switch(config)#    | Enables the control traffic QoS feature.  |

## Displaying QoS Information

The **show qos** command displays the current QoS settings along with a the number of frames marked high priority. The count is only for debugging purposes and cannot be configured (see [Example 20-2](#)).

### Example 20-2 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)
```

## Default FCC and QoS Settings

[Table 20-1](#) lists the default settings for FCC and QoS features:

**Table 20-1 Default FCC and QoS Settings**

| Parameters          | Default   |
|---------------------|-----------|
| FCC protocol        | Disabled. |
| QoS control traffic | Enabled.  |



# CHAPTER 21

## Configuring System Message Logging

---

This chapter describes how to configure system message logging on the Cisco MDS 9000 Family switches. It includes the following sections:

- [About System Message Logging, page 21-2](#)
- [System Log Message Format, page 21-4](#)
- [Configuring System Message Logging, page 21-5](#)
- [Displaying System Message Logging Information, page 21-8](#)
- [Default Settings, page 21-13](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## About System Message Logging

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information.
- Allows you to select the destination of the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility (see [Table 21-1](#)) and the severity level (see [Table 21-2](#)). Messages are time-stamped to enhance real-time debugging and management.

You can access logged system messages using the CLI or by saving them to a properly configured syslog server. The switch software saves syslog messages in a file that can be configured to save up to 4 MB. You can monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a syslog server.



### Note

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a syslog 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. You can view this log at any time using the **show logging nvram** command.

[Table 21-1](#) describes the facilities supported by the system message logs.

**Table 21-1 Internal Logging Facilities**

| Facility Keyword | Description                    | Standard or Cisco MDS Specific |
|------------------|--------------------------------|--------------------------------|
| <b>acl</b>       | ACL manager                    | Cisco MDS 9000 Family specific |
| <b>all</b>       | All facilities                 | Cisco MDS 9000 Family specific |
| <b>auth</b>      | Authorization system           | Standard                       |
| <b>authpriv</b>  | Authorization (private) system | Standard                       |
| <b>bootvar</b>   | Bootvar                        | Cisco MDS 9000 Family specific |
| <b>callhome</b>  | Call Home                      | Cisco MDS 9000 Family specific |
| <b>cron</b>      | Cron or at facility            | Standard                       |
| <b>daemon</b>    | System daemons                 | Standard                       |
| <b>fcc</b>       | FCC                            | Cisco MDS 9000 Family specific |
| <b>fcdomain</b>  | fcdomain                       | Cisco MDS 9000 Family specific |
| <b>fcns</b>      | Name server                    | Cisco MDS 9000 Family specific |
| <b>fcs</b>       | FCS                            | Cisco MDS 9000 Family specific |
| <b>flogi</b>     | FLOGI                          | Cisco MDS 9000 Family specific |
| <b>fspf</b>      | FSPF                           | Cisco MDS 9000 Family specific |
| <b>ftp</b>       | File Transfer Protocol         | Standard                       |



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Table 21-1 Internal Logging Facilities**

| Facility Keyword        | Description                      | Standard or Cisco MDS Specific |
|-------------------------|----------------------------------|--------------------------------|
| <b>ipconf</b>           | IP configuration                 | Cisco MDS 9000 Family specific |
| <b>ipfc</b>             | IPFC                             | Cisco MDS 9000 Family specific |
| <b>kernel</b>           | Kernel                           | Standard                       |
| <b>local0 to local7</b> | Locally defined messages         | Standard                       |
| <b>lpr</b>              | Line printer system              | Standard                       |
| <b>mail</b>             | Mail system                      | Standard                       |
| <b>mcast</b>            | Multicast                        | Cisco MDS 9000 Family specific |
| <b>module</b>           | Switching module                 | Cisco MDS 9000 Family specific |
| <b>news</b>             | USENET news                      | Standard                       |
| <b>ntp</b>              | NTP                              | Cisco MDS 9000 Family specific |
| <b>platform</b>         | Platform manager                 | Cisco MDS 9000 Family specific |
| <b>port</b>             | Port                             | Cisco MDS 9000 Family specific |
| <b>port-channel</b>     | PortChannel                      | Cisco MDS 9000 Family specific |
| <b>qos</b>              | QoS                              | Cisco MDS 9000 Family specific |
| <b>rdl</b>              | RDL                              | Cisco MDS 9000 Family specific |
| <b>rib</b>              | RIB                              | Cisco MDS 9000 Family specific |
| <b>rscn</b>             | RSCN                             | Cisco MDS 9000 Family specific |
| <b>securityd</b>        | Security                         | Cisco MDS 9000 Family specific |
| <b>syslog</b>           | Internal syslog messages         | Standard                       |
| <b>sysmgr</b>           | System manager                   | Cisco MDS 9000 Family specific |
| <b>tlport</b>           | TL port                          | Cisco MDS 9000 Family specific |
| <b>user</b>             | User process                     | Standard                       |
| <b>uucp</b>             | Unix-to-Unix copy system         | Standard                       |
| <b>vhbad</b>            | Virtual host base adapter daemon | Cisco MDS 9000 Family specific |
| <b>vni</b>              | Virtual network interface        | Cisco MDS 9000 Family specific |
| <b>vrrp_cfg</b>         | VRRP configuration               | Cisco MDS 9000 Family specific |
| <b>vrrp_eng</b>         | VRRP engine                      | Cisco MDS 9000 Family specific |
| <b>vsan</b>             | VSAN syslog                      | Cisco MDS 9000 Family specific |
| <b>vshd</b>             | vshd                             | Cisco MDS 9000 Family specific |
| <b>wwn</b>              | WWN manager                      | Cisco MDS 9000 Family specific |
| <b>xbar</b>             | Xbar syslog                      | Cisco MDS 9000 Family specific |
| <b>zone</b>             | Zone server                      | Cisco MDS 9000 Family specific |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

Table 21-2 describes the severity levels supported by the system message logs.

**Table 21-2 Error Message Severity Levels**

| Level Keyword | Level | Description                      | Syslog 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         |

## System Log Message Format

System log messages begin with a percent sign (%) and are displayed in the following format (see Table 21-3):

```
month dd hh:mm:ss switchname facility-severity-MNEMONIC description
```

For example:

```
Nov 8 14:07:58 excal-113 %LOG_MODULE-5-MOD_OK: Module 1 is online
Nov 8 14:07:58 excal-113 %LOG_PORT-3-IF_UNSUPPORTED_TRANSCEIVER: Transceiver for interface
fc1/13 is not supported
Nov 8 14:07:59 excal-113 %LOG_PLATFORM-5-PS_OK: Power supply 1 ok
Nov 8 14:07:53 excal-113 %LOG_DAEMON-5-SYSTEM_MSG: readjusting service shell
Nov 8 15:59:38 excal-113 %LOG_KERN-6-SYSTEM_MSG: utaker: setting queue 1 control pid 1392
(owner 1392)
Nov 8 15:21:44 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(171.71.58.72)
```

**Table 21-3 System Log Message Format Description**

| Element     | Description                                                                   |
|-------------|-------------------------------------------------------------------------------|
| month dd    | The date and month of the error or event.                                     |
| hh:mm:ss    | The time of the error or event.                                               |
| switchname  | The name of the switch                                                        |
| facility    | The facility of the error or event (daemon, kernel, VSHD, or other facility). |
| severity    | Single-digit code from 0 to 7 that indicates the severity of the message.     |
| MNEMONIC    | Text string that uniquely describes the error message.                        |
| description | Text string containing detailed information about the event being reported    |

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

# Configuring System Message Logging

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

## Enabling Message Logging

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# <b>terminal monitor</b>    | Enables logging for a Telnet, or SSH session.<br><b>Note</b> A console session is enabled by default.         |
| Step 2 | switch# <b>terminal no monitor</b> | Disables logging for a Telnet, or SSH session.<br><b>Note</b> A Telnet or SSH session is disabled by default. |

## Configuring 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).

To configure the severity level for a logging facility, follow these steps:

|        | Command                                     | Purpose                                                                                                                                                                   |
|--------|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)## | Enters configuration mode.                                                                                                                                                |
| Step 2 | switch(config)# <b>logging console 3</b>    | Configures console logging at level 3 (error). Logging messages with a severity level of 3 or above will be displayed on the console.                                     |
|        | switch(config)# <b>logging console</b>      | Reverts console logging to the factory set default severity level of 2 (critical). Logging messages with a severity level of 2 or above will be displayed on the console. |



### 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 Line Console Settings”](#) section on page 3-31).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Configuring 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 configure the severity level for a logging facility, follow these steps:

|        | Command                                    | Purpose                                                                                                                                                                       |
|--------|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)# | Enters configuration mode.                                                                                                                                                    |
| Step 2 | switch(config)# <b>logging module 1</b>    | Configures module logging at Level 1 (alerts).                                                                                                                                |
|        | switch(config)# <b>logging module</b>      | Configures module logging for all modules in the switch.                                                                                                                      |
|        | switch(config)# <b>no logging console</b>  | Reverts console logging to the factory set default severity level of 5 (notification). Logging messages with a severity level of 5 or above will be displayed on the console. |

## Configuring Facility Severity Level

To configure the severity level for a logging facility, follow these steps:

|        | Command                                       | Purpose                                                                                                                                                             |
|--------|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#    | Enters configuration mode.                                                                                                                                          |
| Step 2 | switch(config)# <b>logging level kernel 4</b> | 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 will be displayed. |

## Configuring Log Files

Logging messages may be saved to a log file. You can configure the name of this file and restrict its size as required. The file name can have up to 200 characters and the file size ranges from 4096 bytes to 4194304 bytes.

To send log messages to file, follow these steps:

|        | Command                                                                    | Purpose                                                                                                                                                                                                                                      |
|--------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                 | Enters configuration mode.                                                                                                                                                                                                                   |
| Step 2 | switch(config)# <b>logging logfile</b><br><b>ManagerLog 3 size 3000000</b> | Configures logging information for errors or events above severity level 3 to be logged in a file named ManagerLog. By configuring a size, you are restricting the file size to 3000000 bytes. The maximum upper limit is 4194304 (default). |

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** and **clear debug-logfile** commands to view and clear this file. It is not accessible using the **dir** command.

You can display the log file using the **show logging logfile** command and copy the logfile to a different location using the **copy log:messages** command using additional copy syntax (see the “[Copying Files](#)” section on page 3-28).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Configuring Syslog Servers

To send log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. Log in as root, and perform these steps:

**Step 1** Add the following line to the file `/etc/syslog.conf`

```
local7.debug /var/log/myfile.log
```



**Note** Be sure to add five tab characters between **local7.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 **local7** 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 syslog daemon reads the new changes by entering this command:

```
$ kill -HUP -cat /etc/syslog.pid~
```

To configure syslog servers, follow these steps:

|               | Command                                                                       | Purpose                                                                                                                                                                                                                                                                |
|---------------|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b><br>switch(config)#                                    | Enters configuration mode.                                                                                                                                                                                                                                             |
| <b>Step 2</b> | switch(config)# <b>logging server</b><br><b>172.22.00.00</b>                  | 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 IP address (172.22.00.00).<br><br><b>Note</b> You can configure a maximum of three syslog servers. |
|               | switch(config)# <b>logging server</b><br><b>172.22.00.00 facility local 1</b> | Configures the switch to forward log messages according to the specified facility (local1) for the server IP address (172.22.00.00). The default outgoing facility is local7.                                                                                          |
|               | switch(config)# <b>no logging server</b><br><b>172.11.00.00</b>               | Removes the specified server (172.11.00.00) and reverts to factory default.<br><br><b>Note</b> You can configure a maximum of three syslog servers.                                                                                                                    |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Outgoing Syslog Server Logging Facilities

All syslog 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 syslog 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 21-1](#) and the outgoing logging facilities are listed in [Table 21-4](#).

**Table 21-4 Outgoing Logging Facilities**

| Facility Keyword        | Description                    | Standard or Cisco MDS Specific   |
|-------------------------|--------------------------------|----------------------------------|
| <b>auth</b>             | Authorization system           | Standard                         |
| <b>authpriv</b>         | Authorization (private) system | Standard                         |
| <b>cron</b>             | Cron or at facility            | Standard                         |
| <b>daemon</b>           | System daemons                 | Standard                         |
| <b>ftp</b>              | File Transfer Protocol         | Standard                         |
| <b>kernel</b>           | Kernel                         | Standard                         |
| <b>local0 to local7</b> | Locally defined messages       | Standard (local7 is the default) |
| <b>lpr</b>              | Line printer system            | Standard                         |
| <b>mail</b>             | Mail system                    | Standard                         |
| <b>news</b>             | USENET news                    | Standard                         |
| <b>syslog</b>           | Internal syslog messages       | Standard                         |
| <b>user</b>             | User process                   | Standard                         |
| <b>uucp</b>             | Unix-to-Unix copy system       | Standard                         |

## Displaying System Message Logging Information

Use the **show logging** command to display the current system message logging configuration. See Examples [21-1](#) to [21-10](#).

**Example 21-1 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
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

| Facility<br>----- | Default Severity<br>----- | Current Session Severity<br>----- |
|-------------------|---------------------------|-----------------------------------|
| 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
...
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

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 21-2 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 21-3 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 21-4 Displays Console Logging Status**

```
switch# show logging console
Logging console: enabled (Severity: notifications)
```

#### **Example 21-5 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                        |



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

|                |              |                  |
|----------------|--------------|------------------|
| 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 21-6 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                        |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

|                |              |                  |
|----------------|--------------|------------------|
| 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 21-7** 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)
```



#### **Note**

Use the `show logging filename` command to display the entire log file.

#### **Example 21-8** Displays Switching Module Logging Status

```
switch# show logging module
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

Logging linecard: enabled (Severity: debugging)

#### Example 21-9 Displays Monitor Logging Status

```
switch# show logging monitor
Logging monitor: enabled (Severity: information)
```



#### Note

Use the **show logging nvram** command to view the log messages in NVRAM.

#### Example 21-10 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 21-5 lists the default settings for system message logging.

**Table 21-5 Default System Message Log Setting**

| Parameters                                | Default         |
|-------------------------------------------|-----------------|
| System message logging to the console     | Enabled.        |
| System message logging to Telnet sessions | Disabled.       |
| Logging file size                         | 4194304.        |
| Log file name                             | 200 characters. |
| Logging server                            | Disabled.       |
| Syslog server IP address                  | Non configured. |
| No. of servers                            | 3 servers.      |
| Server facility                           | Local 7.        |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***



## 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, page 22-1](#)
- [Starting SCSI LUN Discovery, page 22-2](#)
- [Initiating Customized Discovery, page 22-2](#)
- [Displaying SCSI LUN Information, page 22-2](#)

### 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

To begin SCSI LUN discovery, follow this step:

|        | Command                                                                                                                                                                                                                                                                                                      | Purpose                                                                 |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Step 1 | switch# <b>discover scsi-target local</b><br>discovery started                                                                                                                                                                                                                                               | Discovers local SCSI targets.                                           |
|        | switch# <b>discover scsi-target remote</b><br>discovery started                                                                                                                                                                                                                                              | Discovers remote SCSI targets.                                          |
|        | switch# <b>discover scsi-target vsan 1 fcid 0x9c03d6</b><br>discover scsi-target vsan 1 fcid 0x9c03d6<br>VSAN: 1 FCID: 0x9c03d6 PWWN:<br>00:00:00:00:00:00:00:00<br>PRLI RSP: 0x01 SPARM: 0x0012<br>SCSI TYPE: 0 NLUNS: 1<br>Vendor: Company 4 Model: ST318203FC Rev: 0004<br>Other: 00:00:02:32:8b:00:50:0a | Discovers SCSI targets for the specified VSAN (1) and FC ID (0x9c03d6). |
|        | switch# <b>discover scsi-target custom-list</b><br>discovery started                                                                                                                                                                                                                                         | Discovers SCSI targets from the customized list.                        |

Only Nx ports present in the name server database and which have registered a FC4 Type = SCSI\_FCP are discovered.

## 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.

To initiate a customized discovery, follow this step:

|        | Command                                                                      | Purpose                                               |
|--------|------------------------------------------------------------------------------|-------------------------------------------------------|
| Step 1 | switch# <b>discover custom-list add vsan 1 domain 0X123456</b><br>switch#    | Adds the specified entry to the custom list.          |
|        | switch# <b>discover custom-list delete vsan 1 domain 0X123456</b><br>switch# | 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 22-1 to 22-5.

### Example 22-1 Displays the Discovered Targets

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

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Note**

The discovery can take several minutes to complete, especially if the fabric is large fabric or if several devices are slow to respond.

**Example 22-2 Displays the FCNS Database**

```
switch# show fcns database
VSAN 1:
```

| FCID     | TYPE | PWWN                    | (VENDOR)    | FC4-TYPE:FEATURE |
|----------|------|-------------------------|-------------|------------------|
| 0x9c0000 | N    | 21:00:00:e0:8b:08:96:22 | (Company 1) | scsi-fcp:init    |
| 0x9c0100 | N    | 10:00:00:05:30:00:59:1f | (Company 2) | ipfc             |
| 0x9c0200 | N    | 21:00:00:e0:8b:07:91:36 | (Company 3) | scsi-fcp:init    |
| 0x9c03d6 | NL   | 21:00:00:20:37:46:78:97 | (Company 4) | scsi-fcp:target  |
| 0x9c03d9 | NL   | 21:00:00:20:37:5b:cf:b9 | (Company 4) | scsi-fcp:target  |
| 0x9c03da | NL   | 21:00:00:20:37:18:6f:90 | (Company 4) | scsi-fcp:target  |
| 0x9c03dc | NL   | 21:00:00:20:37:5a:5b:27 | (Company 4) | scsi-fcp:target  |
| 0x9c03e0 | NL   | 21:00:00:20:37:36:0b:4d | (Company 4) | scsi-fcp:target  |
| 0x9c03e1 | NL   | 21:00:00:20:37:39:90:6a | (Company 4) | scsi-fcp:target  |
| 0x9c03e2 | NL   | 21:00:00:20:37:18:d2:45 | (Company 4) | scsi-fcp:target  |
| 0x9c03e4 | NL   | 21:00:00:20:37:6b:d7:18 | (Company 4) | scsi-fcp:target  |
| 0x9c03e8 | NL   | 21:00:00:20:37:38:a7:c1 | (Company 4) | scsi-fcp:target  |
| 0x9c03ef | NL   | 21:00:00:20:37:18:17:d2 | (Company 4) | scsi-fcp:target  |

Total number of entries = 13

**Example 22-3 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 22-4 Displays the Discovered LUNs**

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

```
- ST318203FC from Company 4 (Rev 0004)
```

```
FCID is 0x9c03d6 in VSAN 1, PWWN is 21:00:00:20:37:46:78:97
```

| LUN | Capacity<br>(MB) | Status | Serial Number | Device-Id |
|-----|------------------|--------|---------------|-----------|
|-----|------------------|--------|---------------|-----------|

|     |       |        |                  |                                     |
|-----|-------|--------|------------------|-------------------------------------|
| 0x0 | 18210 | Online | LRA2510000007027 | C:1 A:0 T:3 20:00:00:20:37:46:78:97 |
|-----|-------|--------|------------------|-------------------------------------|

```
- ST318203FC from Company 4 (Rev 0004)
```

```
FCID is 0x9c03d9 in VSAN 1, PWWN is 21:00:00:20:37:5b:cf:b9
```

| LUN | Capacity<br>(MB) | Status | Serial Number | Device-Id |
|-----|------------------|--------|---------------|-----------|
|-----|------------------|--------|---------------|-----------|

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```

0x0 18210 Online LR94873000007029 C:1 A:0 T:3 20:00:00:20:37:5b:cf:b9
- ST318203FC from Company 4 (Rev 0004)
FCID is 0x9c03da in VSAN 1, PWWN is 21:00:00:20:37:18:6f:90

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18210 Online LR18591800001004 C:1 A:0 T:3 20:00:00:20:37:18:6f:90
- ST318203FC from Company 4 (Rev 0004)
FCID is 0x9c03dc in VSAN 1, PWWN is 21:00:00:20:37:5a:5b:27

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18210 Online LRC4498200007031 C:1 A:0 T:3 20:00:00:20:37:5a:5b:27
- ST318203FC from Company 4 (Rev 0004)
FCID is 0x9c03e0 in VSAN 1, PWWN is 21:00:00:20:37:36:0b:4d

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18210 Online LR18184700007024 C:1 A:0 T:3 20:00:00:20:37:36:0b:4d
- ST318203 CLAR18 from Company 4 (Rev 3844)
FCID is 0x9c03e1 in VSAN 1, PWWN is 21:00:00:20:37:39:90:6a

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18200 Online LR64147100001017 C:1 A:0 T:3 20:00:00:20:37:39:90:6a
- ST318203 CLAR18 from Company 2 (Rev 3844)
FCID is 0x9c03e2 in VSAN 1, PWWN is 21:00:00:20:37:18:d2:45

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18200 Online LR28349500001952 C:1 A:0 T:3 20:00:00:20:37:18:d2:45
- ST318203 CLAR18 from Company 2 (Rev 3844)
FCID is 0x9c03e4 in VSAN 1, PWWN is 21:00:00:20:37:6b:d7:18

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18200 Online LRF7150500001041 C:1 A:0 T:3 20:00:00:20:37:6b:d7:18
- ST318203FC from Company 2 (Rev 0004)
FCID is 0x9c03e8 in VSAN 1, PWWN is 21:00:00:20:37:38:a7:c1

LUN Capacity Status Serial Number Device-Id
 (MB)

0x0 18210 Online LR43588300001011 C:1 A:0 T:3 20:00:00:20:37:38:a7:c1
...

```

**Example 22-5 Displays Customized Discovered Targets**

```

switch# show scsi-target custom-list

VSAN DOMAIN

1 56

```





## 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, page 23-2](#)
- [SPAN Sources, page 23-2](#)
- [SPAN Sessions, page 23-5](#)
- [Specifying Filters, page 23-5](#)
- [SD Port Characteristics, page 23-5](#)
- [Configuring SPAN, page 23-6](#)
- [Monitoring Traffic Using Fibre Channel Analyzers, page 23-10](#)
- [Displaying SPAN Information, page 23-13](#)
- [Default Settings, page 23-14](#)


**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## About SPAN

The switched port analyzer (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 “Configuring a Fabric Analyzer” section on page 24-5).

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 23-1).

**Figure 23-1 SPAN Transmission**



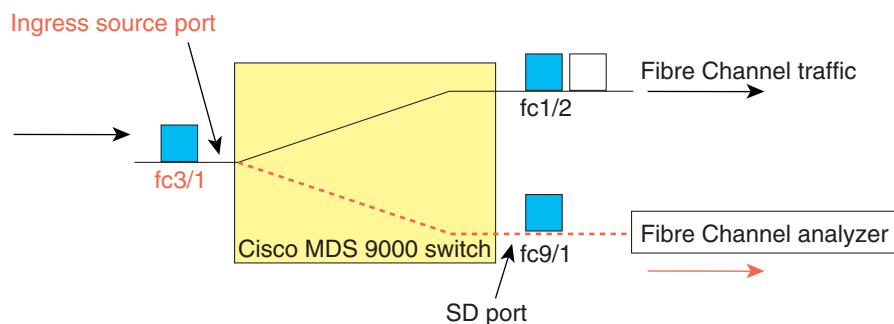
85478

## 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 23-2).

**Figure 23-2 SPAN Traffic from the Ingress Direction**




85479

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

- Egress source (tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 23-3).

**Figure 23-3 SPAN Traffic from Egress Direction**



## Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports:
  - 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.

## VSAN as a SPAN 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.

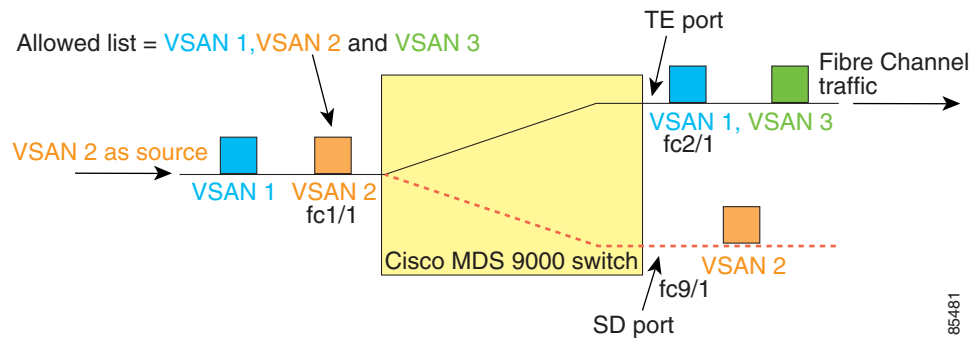
**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## 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.
- When a VSAN is specified as a source, you will not be able to 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 SPAN source, you will not be able to 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 23-4](#) displays a configuration using VSAN 2 as a SPAN 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 23-4 VSAN As a SPAN Source**



For this configuration, the following apply:

- VSAN 2 as a SPAN 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 as the port VSAN does not match VSAN 1. See “[Configuring Trunk-Allowed VSAN List](#)” section on page 10-4 or “[VSAN Membership](#)” section on page 8-6.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## 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 a SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic will not be directed to the SD port.

To temporarily deactivate (suspend) a SPAN session use the **suspend** command in the SPAN submode. The traffic monitoring is stopped during this time. You can reactivate the SPAN session using the **no suspend** command.



### Tip

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

## 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 23-4](#)). Only VSANs present in the filter are spanned.

You can specify session VSAN filters which 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 buffer-to-buffer 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 port mode can not be changed if it is being used for a SPAN session.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*



**Note**

If you need to change a 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.

- The outgoing frames can be encapsulated in extended inter-switch link (EISL) format.
- The SD port does not have a port VSAN.

# Guidelines to Configure SPAN

The following guidelines apply for a SPAN configuration:

- 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 [“Configuring 32-port Switching Modules” section on page 9-7](#)).
- 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 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                                          |
|---------------|-------------------------------------------------|--------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b>                         | Enters configuration mode.                       |
| <b>Step 2</b> | switch(config)# <b>interface fc9/1</b>          | Configures the specified interface.              |
| <b>Step 3</b> | switch(config-if)# <b>switchport mode SD</b>    | Configures the SD port-mode for interface fc2/1. |
| <b>Step 4</b> | switch(config-if)# <b>switchport speed 1000</b> | Configures the SD port speed to 1000 Mbps.       |
| <b>Step 5</b> | switch(config-if)# <b>no shutdown</b>           | Enables traffic flow.                            |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

To configure a SPAN session, follow these steps:

|        | Command                                                          | Purpose                                                                                       |
|--------|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                          | Enters configuration mode.                                                                    |
| Step 2 | switch(config)# <b>span session 1</b><br>switch(config-span)#    | Configures the specified SPAN session (1). If the session does not exist, it will be created. |
|        | switch(config)# <b>no span session 1</b>                         | Deletes the specified SPAN session (1).                                                       |
| Step 3 | switch(config-span)# <b>destination interface fc9/1</b>          | Configures the specified destination interface (fc 9/1) in a session.                         |
|        | switch(config-span)# <b>no destination interface fc9/1</b>       | Removes the specified destination interface (fc 9/1).                                         |
| Step 4 | switch(config-span)# <b>source interface fc7/1</b>               | Configures the source (fc7/1) interface in both directions.                                   |
|        | switch(config-span)# <b>no source interface fc7/1</b>            | Removes the specified destination interface (fc 7/1) from this session.                       |
| Step 5 | switch(config-span)# <b>source interface sup-fc0</b>             | Configures the source interface (sup-fc0) in the session.                                     |
|        | switch(config-span)# <b>source interface fc1/5 - 6, fc2/1 -3</b> | Configures the specified interface ranges in the session.                                     |
|        | switch(config-span)# <b>source vsan 1-2</b>                      | Configures source VSANs 1 and 2 in the session.                                               |
|        | switch(config-span)# <b>source interface port-channel 1</b>      | Configures the source PortChannel (port-channel 1).                                           |
|        | switch(config-span)# <b>no source interface port-channel 1</b>   | Deletes the specified source interface (port-channel 1)                                       |
| Step 6 | switch(config-span)# <b>suspend</b>                              | Suspends the session.                                                                         |
|        | switch(config-span)# <b>no suspend</b>                           | Reactivates the session.                                                                      |

To configure a SPAN filter, follow these steps:

|        | Command                                                       | Purpose                                                              |
|--------|---------------------------------------------------------------|----------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                       | Enters configuration mode.                                           |
| Step 2 | switch(config)# <b>span session 1</b><br>switch(config-span)# | Configures the specified session (1).                                |
| Step 3 | switch(config-span)# <b>source interface fc9/1 tx</b>         | Configures the source fc9/1 interface in the egress (tx) direction   |
|        | switch(config-span)# <b>source filter vsan 1-2</b>            | Configures VSANs 1 and 2 as session filters.                         |
|        | switch(config-span)# <b>source interface fc7/1 rx</b>         | Configures the source fc7/1 interface in the ingress (rx) direction. |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Encapsulating Frames

The **switchport encap eisl** command only applies to SD port interfaces. This command is disabled by default. If you enable the encapsulation feature, all outgoing frames will be encapsulated. If encapsulation is enabled, you will see a new line (Encapsulation is eisl) in the **show int SD\_port\_interface** command output.

To encapsulate outgoing frames (optional), follow these steps:

|        | Command                                            | Purpose                                                                      |
|--------|----------------------------------------------------|------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                            | Enters configuration mode.                                                   |
| Step 2 | switch(config)# <b>interface fc9/32</b>            | Configures the specified interface.                                          |
| Step 3 | switch(config-if)# <b>switchport mode SD</b>       | Configures the SD port-mode for interface fc2/1.                             |
| Step 4 | switch(config-if)# <b>switchport encap eisl</b>    | Enables the encapsulation option for this SD port.                           |
| Step 5 | switch(config-if)# <b>no switchport encap eisl</b> | Disables the encapsulation option and reverts the switch to factory default. |

## SPAN Conversion Behavior

Effective Release 1.1(1), SPAN features (configured in any prior release) are converted as stated below:

- 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 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 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 will be removed from both directions. For example,

Before Release 1.0(4):

```
Session 2 (active)
 Destination is fc1/9
 No session filters configured
 Ingress (rx) sources are
```



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```
vsans 12
fc1/6 (vsan 1-20),
Egress (tx) sources are
fc1/6 (vsan 1-20),
```

Once upgraded to 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
```

Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Release 1.0(4). When upgraded to 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).

**Note**

---

The deprecated configurations are removed from persistent memory, once a switchover or a new startup configuration is implemented.

---

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## Monitoring Traffic Using Fibre Channel Analyzers

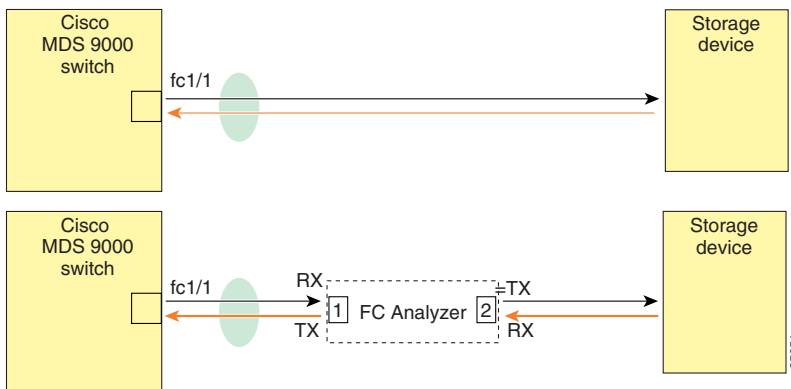
You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is specially useful in troubleshooting scenarios when traffic disruption changes the problem environment and makes it difficult to reproduce the problem.

### 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 23-5](#).

**Figure 23-5 Fibre Channel Analyzer Usage Without SPAN**

FC Analyzer usage without SPAN



This type of connection has the following limitations:

- 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.

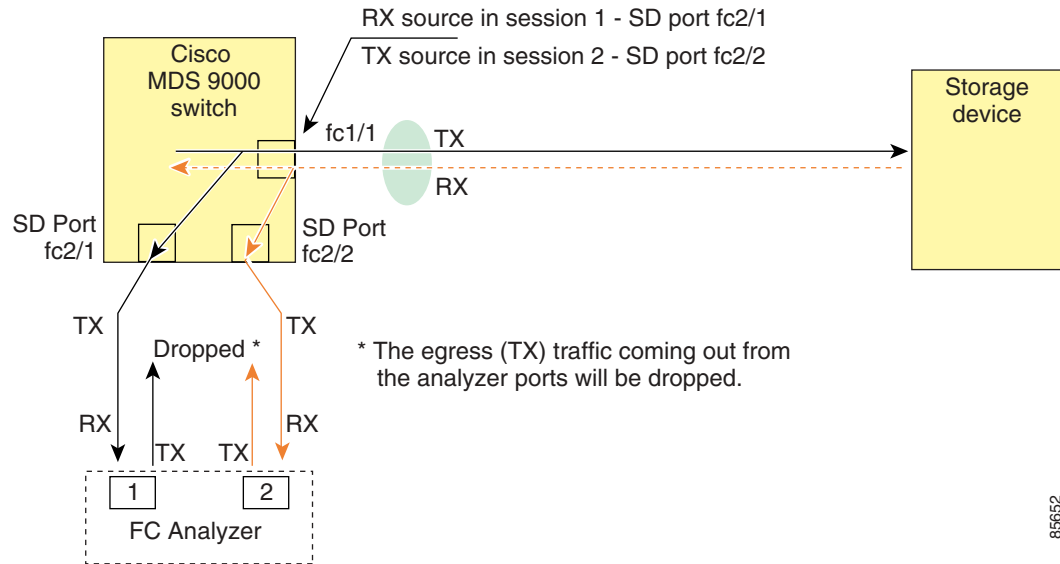
### Using SPAN

Using SPAN you can capture the same traffic scenario shown in [Figure 23-5](#) 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 23-6](#).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Figure 23-6 Fibre Channel Analyzer Using SPAN**



85652

## Configuring Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in [Figure 23-6](#), 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 fc1/1 in 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                                                         |
|---------------|---------------------------------------------------------------|-----------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b>                                       | Enters configuration mode.                                      |
| <b>Step 2</b> | switch(config)# <b>span session 1</b><br>switch(config-span)# | Creates the SPAN session 1.                                     |
| <b>Step 3</b> | switch(config-span)## <b>destination interface fc2/1</b>      | Configures the destination interface fc2/1.                     |
| <b>Step 4</b> | switch(config-span)# <b>source interface fc1/1 rx</b>         | Configures the source interface fc1/1 in the ingress direction. |
| <b>Step 5</b> | switch(config)# <b>span session 2</b><br>switch(config-span)# | Creates the SPAN session 2.                                     |
| <b>Step 6</b> | switch(config-span)## <b>destination interface fc2/2</b>      | Configures the destination interface fc2/2.                     |
| <b>Step 7</b> | switch(config-span)# <b>source interface fc1/1 tx</b>         | Configures the source interface fc1/1 in the egress direction.  |

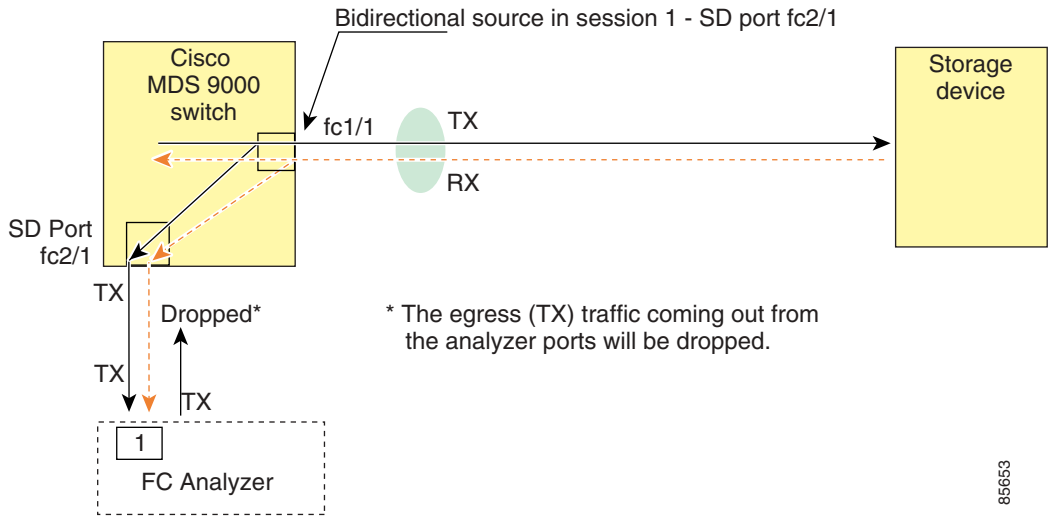
[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Using a 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 23-6. 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 23-7 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 direction. This setup is more advantageous and cost-effective than the setup shown in Figure 23-6—it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 23-7 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# <b>config t</b>                                       | Enters configuration mode.                                 |
| Step 2 | switch(config)# <b>span session 1</b><br>switch(config-span)# | Creates the SPAN session 1.                                |
| Step 3 | switch(config-span)## <b>destination interface fc2/1</b>      | Configures the destination interface fc2/1.                |
| Step 4 | switch(config-span)# <b>source interface fc1/1</b>            | Configures the source interface fc1/1 on the same SD port. |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Displaying SPAN Information

Use the **show span** command to display configured SPAN information. See Examples 23-1 to 23-4.

### Example 23-1 Displays SPAN Sessions in a Brief Format

```
switch# show span session brief
```

| Session | Admin State | Oper State | Destination Interface |
|---------|-------------|------------|-----------------------|
| 7       | no suspend  | active     | fc2/7                 |
| 1       | suspend     | inactive   | not configured        |
| 2       | no suspend  | inactive   | fc3/1                 |

### Example 23-2 Displays a Specific SPAN Session Details

```
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 23-3 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,
 port-channel 2, sup-fc0,
 Egress (tx) sources are
 fc3/2, fc3/3, fc3/4, sup-fc0,
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Example 23-4** 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
```

# Default Settings

Table 23-1 lists the default settings for SPAN parameters

**Table 23-1** 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.                                             |



## Advanced Features and Concepts

---

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [Configuring Time Out Values, page 24-2](#)
- [Invoking the fctrace Feature, page 24-3](#)
- [Invoking the fcping Feature, page 24-4](#)
- [Configuring a Fabric Analyzer, page 24-5](#)
- [Configuring World Wide Names, page 24-16](#)
- [Allocating Flat FC IDs, page 24-18](#)
- [Enabling Loop Monitoring, page 24-18](#)
- [Configuring the Switch for Interoperability, page 24-20](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## Configuring Time Out Values

The **ftimer** command modifies Fibre Channel protocol related timer values for the switch. You can only configure Fibre Channel time out values (TOVs) commands if all VSANs in a switch are suspended.



### Note

The F\_S\_TOV constant can not be configured.

You can use the **ftimer** command in configuration mode to configure the following 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.
- Resource allocation TOV (R\_A\_TOV)—the valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds.



### Caution

These values can not be changed unless all VSANs in the switch are suspended.

If you issue the **ftimer** command without suspending all VSANs in a switch, you will get a warning message:

```
switch(config)# ftimer D_S_TOV 6000
Warning: This configuration would impact whole fabric.
Since this configuration is not propagated to other switches.
Please configure the same value in all the switches
It is recommended that all vsans be suspended before executing this command
suspend all vsans first
could not update the value
```

Use the **show ftimer** command to display show the configured ftimer values (see [Example 24-1](#)).

### Example 24-1 Displays Configured TOVs

```
switch# show ftimer
F_S_TOV : 5000 milliseconds
D_S_TOV : 5000 milliseconds
E_D_TOV : 2000 milliseconds
R_A_TOV : 10000 milliseconds
```



### Note

The F\_S\_TOV constant, though not configured, is displayed in the output of the **show ftimer** command.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Invoking the fctrace Feature

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 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.

To perform a fctrace operation, follow this step:

|        | Command                                                                                                                                                                                                                                                             | Purpose                                                                                                                                                                        |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <pre>switch# fctrace fcid 0xd70000 vsan 1 Route present for : 0xd70000 20:00:00:0b:46:00:02:82(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre>                                         | Invokes fctrace for the specified FC ID of the destination N port                                                                                                              |
|        | <pre>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(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre> | <p>Invokes fctrace using the pWWN of the destination N port</p> <p>By default the period to wait before timing out is 5 seconds, The range is from one through 10 seconds.</p> |



### Note

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.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Invoking the fcping Feature

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 or the destination port WWN information.

To perform a fcping operation, follow this step:

|        | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Purpose                                                                                                                                                                                                   |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <pre>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</pre>                                                                                                                                                                                                               | Performs a fcping operation for the specified pWWN or the FCID of the destination. By default, five frames are sent.                                                                                      |
|        | <pre>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</pre> | Sets the number of frames to be sent using the count option. The range is from 0 through 2147483647. A value of 0 will ping forever.                                                                      |
|        | <pre>switch# fcping fcid 0xd500b4 vsan 1 timeout 10 28 bytes from 0xd500b4 time = 1345 usec 28 bytes from 0xd500b4 time = 417 usec 28 bytes from 0xd500b4 time = 340 usec 28 bytes from 0xd500b4 time = 451 usec 28 bytes from 0xd500b4 time = 356 usec  5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 340/581/1345 usec</pre>                                                                                                                                                                                                  | Sets the timeout value. The default period to wait is 5 seconds. The range is from 1 through 10 seconds.                                                                                                  |
| Step 2 | <pre>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 28 bytes from 21:00:00:20:37:6f:db:dd time = 471 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 372 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 364 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 1261 usec  5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 364/784/1454 usec</pre>                                                | <p>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.</p> <p>Retry the command a few seconds later.</p> |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. While 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.

Cisco has brought protocol analysis within a storage network to a new capability level 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.

Cisco's Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—You can obtain more information from <http://www.tcpdump.org>.
- Ethereal—You can obtain more information from <http://www.ethereal.com>.



### Note

Cisco's 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 explains the following topics:

- [About the Cisco Fabric Analyzer, page 24-5](#)
- [Configuring the Cisco Fabric Analyzer, page 24-7](#)
- [Viewing Display Filters Information, page 24-10](#)
- [Clearing Configured fcanalyzer Information, page 24-9](#)
- [Display Filters, page 24-10](#)


## About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer comprises two separate components (see [Figure 24-1](#)):

- A 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
- A 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.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Figure 24-1 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 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 switch, it is protected by the roles-based policy that limits access in each switch.

See the [“Capturing Frames Locally”](#) section on page 24-7.

## 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 end points, 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 unauthorized machine in the network from snooping on the control traffic in the network.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

RPCAP supports two setup connection modes based on fire wall 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 24-9.

## 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>. Since Ethereal has a GUI front-end, it supports a rich functionality such as colorized display, graphical assists in defining filters, and searching for specific frames. These features are documented on Ethereal’s web site.

While remote capture via 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 “Display Filters” section on page 24-10.

## Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer by issuing the **fcanalyzer local** or **fcanalyzer remote** commands in configuration mode.

- **Local capture**—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby.
- **Remote capture**—The command setting to enable a remote capture can be saved to persistent storage using the **copy** command. 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

Launches the textual version on the analyzer directly on the console screen. The capture can also be saved on the local file system.

To capture frames locally, follow these steps:

|        | Command                                                              | Purpose                    |
|--------|----------------------------------------------------------------------|----------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                           | Enters configuration mode. |
|        | <b>Note</b> The options within Step 2 may be performed in any order. |                            |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

|        | Command                                                                                                                                                                                                                                                                                         | Purpose                                                                                                                                                                                 |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 2 | switch(config)# <b>fcanalyzer local</b><br>Capturing on eth2<br>switch(config)#                                                                                                                                                                                                                 | Begins capturing the frames locally (supervisor module).                                                                                                                                |
|        | switch(config)# <b>fcanalyzer local brief</b><br>Capturing on eth2<br>switch(config)#                                                                                                                                                                                                           | Displays the protocol summary in a brief format.                                                                                                                                        |
|        | switch(config)# <b>fcanalyzer local display-filter SampleF</b><br>Capturing on eth2                                                                                                                                                                                                             | Displays the filtered frames.                                                                                                                                                           |
|        | switch(config)# <b>fcanalyzer local limit-frame-size 64</b><br>Capturing on eth2<br>switch(config)#                                                                                                                                                                                             | Limits the size of the frame capture to the first 64 bytes. The allowed range is 64 to 65536 bytes.                                                                                     |
|        | switch(config)# <b>fcanalyzer local limit-captured-frames 10</b><br>Capturing on eth2<br>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. |
|        | <b>Note</b> Press <b>Ctrl-c</b> to stop a capture. Otherwise, the capture stops automatically after capturing 100 frames. You can change this default using the <b>fcanalyzer local limit-captured-frames <i>number</i></b> command.                                                            |                                                                                                                                                                                         |
| Step 3 | switch(config)# <b>fcanalyzer local write volatile:sample</b><br>Capturing on eth2<br>switch(config)#                                                                                                                                                                                           | Saves the captured frames to a specified file (sample) in the volatile: directory.<br><br><b>Note</b> Optionally, you can save the specified file to the slot0: directory.              |
|        | <b>Note</b> The final filename that is the capture file will be 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 10MB. |                                                                                                                                                                                         |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## 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, follow these steps:

|        | Command                                                                        | Purpose                                                                                                                                                                                                                                 |
|--------|--------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                     | Enters configuration mode.                                                                                                                                                                                                              |
| Step 2 | switch(config)# <b>fcanalyzer remote 10.21.0.3</b><br>switch(config)#          | Configures the remote IP address (10.21.0.3) to which the captured frames will be sent.                                                                                                                                                 |
|        | switch(config)# <b>fcanalyzer remote 10.21.0.3 active</b><br>switch(config)#   | Enables active mode (passive is the default) with the remote host.<br><br>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)# <b>fcanalyzer remote 10.21.0.3 active 1</b><br>switch(config)# | 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:

- To specify the capture interface 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 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.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Viewing Display Filters Information

Use the **show fcanalyzer** command to display the list of hosts configured for a remote capture. See [Example 24-2](#).

### **Example 24-2 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.

---

## Display Filters

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 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 document in the Ethereal web site (<http://www.ethereal.com>). Some examples of how you can use this feature are provided below:

- To view all packets in a specified VSAN, use this expression:

```
mdshdr.vsan == 2
```

- To view all SW\_ILS frames, use this expression:

```
fcswwils
```

- To view class F frames, use this expression:

```
mdshdr.sof == SOFf
```

- To view all FSPF frames, use this expression:

```
swils.opcode == JLO || 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
```



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## 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>.

## Displaying Filters Examples

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 24-3](#).

### Example 24-3 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)or(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, to observe RSCNs from Fabric Controller and registration and/or query requests to the Name Server. See [Example 24-4](#).



**Note**

The filter requires prior knowledge of the FC ID that is assigned to the N port. Issue the **show flogi database interface** command before running fcanalyzer to obtain the FC ID. In this example, the N port FC ID is 79.03.00.

### Example 24-4 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)or(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)
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

The VSAN ID is specified in hex. See [Example 24-5](#).

**Example 24-5 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 24-6](#).

**Example 24-6 Displays All VSAN 1 Traffic Excluding FSPF hellos and ACK1 Frames.**

```
switch(config)# fcan lo bri dis
(mdshdr.vsan==0x01)&¬((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 24-7](#).

**Example 24-7 Displays SW\_ILS Traffic between Fabric Controllers for all VSANs and Exclude FSPF hellos and ACK1 frames.**

```
switch(config)# fcan lo bri dis
((fc.s_id==ff.ff.fd)&&(fc.type==0x22))&¬(swils.opcode==0x14)
Capturing on eth2
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```

20.573225 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200c 0xffff 0xe -> 0xf ELP
20.574021 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200c 0xacc4 0xff -> 0x0 SW_ACC (ELP)
20.606020 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200d 0xffff 0xe -> 0xf ESC
20.606232 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200d 0xacc5 0xff -> 0x0 SW_ACC (ESC)
20.606665 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200e 0xffff 0xe -> 0xf 0x71
20.608768 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200e 0xacc6 0xff -> 0x0 SW_ACC (0x71)
20.615346 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc7 0xffff 0xff -> 0x0 0x71
20.620330 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc7 0x200f 0xe -> 0xf SW_ACC (0x71)
20.623028 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2010 0xffff 0xe -> 0xf EFP
20.624681 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc9 0xffff 0xff -> 0x0 EFP
20.624974 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2010 0xacc8 0xff -> 0x0 SW_ACC (EFP)
20.625133 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1939 0xffff 0xff -> 0x0 EFP
20.626393 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc9 0x2011 0xe -> 0xf SW_ACC (EFP)
20.627185 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0b 0xffff 0xe -> 0xf EFP
20.627479 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1939 0xab0a 0xe -> 0xf SW_ACC (EFP)
20.627773 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0c 0xffff 0xe -> 0xf DIA
20.631106 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0b 0x193a 0xff -> 0x0 SW_ACC (EFP)
20.631432 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0d 0xffff 0xe -> 0xf MR
20.631567 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193c 0xffff 0xff -> 0x0 DIA
20.631974 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0c 0x193b 0xff -> 0x0 SW_ACC (DIA)
20.631938 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193c 0xab0e 0xe -> 0xf SW_ACC (DIA)
20.639262 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193e 0xffff 0xff -> 0x0 MR
20.640417 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193e 0xab0f 0xe -> 0xf SW_ACC (MR)
20.640598 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0d 0x193d 0xff -> 0x0 SW_ACC (MR)
20.646950 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab14 0xffff 0xe -> 0xf LSU
20.647256 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1944 0xffff 0xff -> 0x0 LSU
20.647996 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1945 0xffff 0xff -> 0x0 LSU
20.648367 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1946 0xffff 0xff -> 0x0 LSA
20.648476 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab17 0xffff 0xe -> 0xf LSU
20.648916 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab19 0xffff 0xe -> 0xf LSA
20.649210 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1a 0xffff 0xe -> 0xf LSA
20.659781 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194a 0xffff 0xff -> 0x0 LSA
20.660535 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1d 0xffff 0xe -> 0xf LSU
20.660649 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194c 0xffff 0xff -> 0x0 LSU
20.660683 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1e 0xffff 0x5 -> 0xf LSU
20.661006 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194e 0xffff 0xff -> 0x0 LSU
20.664994 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab22 0xffff 0xe -> 0xf LSA
20.665341 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab24 0xffff 0x5 -> 0xf LSU
20.665645 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab25 0xffff 0x5 -> 0xf LSA
20.666115 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1952 0xffff 0xff -> 0x0 LSA
20.666445 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1953 0xffff 0xff -> 0x0 LSU
20.666994 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1954 0xffff 0xff -> 0x0 LSA
20.667423 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab2a 0xffff 0x5 -> 0xf LSA
20.667715 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1956 0xffff 0xff -> 0x0 LSA
30.525363 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2012 0xffff 0xe -> 0xf DIA
30.525596 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2012 0xacca 0xff -> 0x0 SW_ACC (DIA)
30.525959 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xaccb 0xffff 0xff -> 0x0 RDI
30.526736 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xaccb 0x2013 0xe -> 0xf SW_ACC (RDI)
30.527032 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2014 0xffff 0xe -> 0xf EFP
30.527662 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2014 0xacc 0xff -> 0x0 SW_ACC (EFP)
30.533157 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2015 0xffff 0xe -> 0xf MR
30.534159 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacce 0xffff 0xff -> 0x0 MR
30.534440 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2015 0xaccd 0xff -> 0x0 SW_ACC (MR)
30.534791 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacce 0x2016 0xe -> 0xf SW_ACC (MR)
30.540883 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x201b 0xffff 0xe -> 0xf LSU
30.541068 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd4 0xffff 0xff -> 0x0 LSU
30.541704 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd5 0xffff 0xff -> 0x0 LSA
30.541981 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd6 0xffff 0xff -> 0x0 LSU
30.542087 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x201e 0xffff 0xe -> 0xf LSA
30.542381 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2020 0xffff 0xe -> 0xf LSU
30.542675 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2021 0xffff 0xe -> 0xf LSU
30.542969 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2022 0xffff 0xe -> 0xf LSA
30.543226 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdb 0xffff 0xff -> 0x0 LSU
30.543614 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2024 0xffff 0xe -> 0xf LSA

```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```

30.543751 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdd 0xffff 0xff -> 0x0 LSA
30.544004 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacde 0xffff 0xff -> 0x0 LSA
30.544522 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdf 0xffff 0xff -> 0x0 LSU
30.544553 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2027 0xffff 0xe -> 0xf LSU
30.550961 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xace7 0xffff 0xff -> 0x0 LSA
30.550988 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x202f 0xffff 0xe -> 0xf LSA

```

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 24-8](#).

#### **Example 24-8 Display SW\_ILS traffic between Fabric Domain Controllers for VSAN 1**

```

switch(config)# fcan lo bri dis
mdshdr.vsan==0x01&&(fc.type==0x22)&&((fc.d_id==ff.fc.79)or(fc.s_id==ff.fc.79))
Capturing on eth2
64.053927 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e15 0xffff 0xff -> 0x0 ACA
64.053995 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e16 0xffff 0xff -> 0x0 ACA
64.054599 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e16 0xb1e2 0x5 -> 0xf SW_ACC (ACA)
64.054747 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e15 0x3037 0x4 -> 0xf SW_ACC (ACA)
64.057643 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e17 0xffff 0xff -> 0x0 SFC
64.057696 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e18 0xffff 0xff -> 0x0 SFC
64.058788 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e17 0x3038 0x5 -> 0xf SW_ACC (SFC)
64.059288 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e18 0xb1e3 0x5 -> 0xf SW_ACC (SFC)
64.062011 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e19 0xffff 0xff -> 0x0 UFC
64.062060 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e1a 0xffff 0xff -> 0x0 UFC
64.073513 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e19 0x3039 0x5 -> 0xf SW_ACC (UFC)
64.765306 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e1a 0xb1e4 0x5 -> 0xf SW_ACC (UFC)
64.765572 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e1b 0xffff 0xff -> 0x0 RCA
64.765626 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e1c 0xffff 0xff -> 0x0 RCA
64.766386 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e1b 0x303a 0x4 -> 0xf SW_ACC (RCA)
64.766392 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e1c 0xb1e5 0x5 -> 0xf SW_ACC (RCA)

```

## 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 is useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restricts a capture to the specified frames. No other frames will be visible until you specify a completely new capture.

The syntax for capture filter 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 document in the Ethereal web site (<http://www.ethereal.com>). Some examples of how you can use this feature are provided below:

- 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
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- 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

- 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)

**[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

## Configuring 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. This WWN is independent of other WWNs on each switch. This centralized control of WWN has the following advantages:

- Efficient sharing of WWN space
- Centralized support across switches

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see [Table 24-1](#)).

**Table 24-1 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 worldwide names are only performed as required. They should not be changed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

## Configuring a Secondary MAC Address

To allocate secondary MAC addresses, follow these steps:

|        | Command                                                                                                                                                                                                                                                                                                                                                                                             | Purpose                                                              |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                                                                                                                                                                                                                                                                                                                          | Enters configuration mode.                                           |
| Step 2 | switch(config)# <b>wwn secondary-mac 00:99:55:77:55:55 range 64</b><br>This command CANNOT be undone.<br>Please enter the BASE MAC ADDRESS again: <b>00:99:55:77:55:55</b><br>Please enter the mac address RANGE again: <b>64</b><br>From now on WWN allocation would be based on new MACs.<br>Are you sure? (yes/no) <b>no</b><br>You entered: no. Secondary MAC NOT programmed<br>switch(config)# | Configures the secondary MAC address. This command cannot be undone. |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Displaying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. See Examples 24-9 to 24-12.

### ***Example 24-9 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 24-10 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 24-11 Displays the WWN for a Specific Switch***

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

### ***Example 24-12 Displays the WWN for a Specified VSAN***

```
switch# show wwn vsan 1
VSAN WWN of VSAN# 1 is 20:01:ac:16:5e:52:00:01
```

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## Allocating Flat FC IDs

Fibre Channel standards require a unique FC ID to be allocated to a 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.

Based on the assigned FC ID, some HBAs assume that no other ports have the same area bits and domain. When a target is assigned with a FC ID that has the same area bits, but different port bits, the HBA fails to discover these targets. To isolate these HBAs in a separate area, switches in the Cisco MDS 9000 Family follow a different FC ID allocation scheme. By default, the FC ID allocation mode is **auto**. In the **auto** mode, only HBAs without interop issues are assigned FCIDs with specific ports bits. All other HBAs are assigned FC IDs with a whole area (port bits set to 0).

The three options to allocate FCID are auto (default), none, and flat.

To allocate flat FC IDs, follow these steps:

|        | Command                                                                  | Purpose                                                                                                                                                      |
|--------|--------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                               | Enters configuration mode.                                                                                                                                   |
| Step 2 | switch(config)# <b>fcinterop fcid-allocation none</b><br>switch(config)# | Allocates one area to the N port attached to an F port.                                                                                                      |
|        | switch(config)# <b>fcinterop fcid-allocation flat</b><br>switch(config)# | Allocates a single FC ID to the N port. This option is generally used to conserve FC ID usage.                                                               |
|        | switch(config)# <b>fcinterop fcid-allocation auto</b><br>switch(config)# | Intelligently assigns flat FC ID to N ports which can interoperate in <b>flat</b> mode, otherwise assigns full area to all other ports. This is the default. |



### Caution

Changes to FC IDs are only performed as required. They should not be changed on a daily basis. Changes, if any, should be made by an administrator or individual who is completely familiar with switch operations.

## Enabling Loop Monitoring

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) using the **fcinterop loop-monitor** command. This command enables loop polling for FL ports in a Cisco MDS 9000 Family switch. By default, the **fcinterop loop-monitor** command is disabled.

To enable the loop monitoring feature, follow these steps:

|        | Command                                          | Purpose                                                                                        |
|--------|--------------------------------------------------|------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#       | Enters configuration mode.                                                                     |
| Step 2 | switch(config)# <b>fcinterop loop-monitor</b>    | Enables the loop monitoring feature.                                                           |
|        | switch(config)# <b>no fcinterop loop-monitor</b> | Disables (default) the loop monitoring feature and reverts the switch to the factory defaults. |



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Caution**

Changes to the loop monitoring feature are only performed as required. They should not be changed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## Configuring the Switch for Interoperability

Interoperability enables multiple vendors' products 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 provide the product with a more aimiable standards compliant implementation.

Table 24-2 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 24-2 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.<br><br>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 setup statically (the MDS will only accept one domain ID, if it doesn't get that domain ID it isolates itself from the fabric), or preferred. (If it doesn't 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                   | 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.                                                                                                                                                                                                                                                                                                                                      |
| Zone propagation          | Some vendors do not pass the full zone configuration ( <b>zoneset</b> ) to other switches, only the active zoneset gets passed.<br><br>Verify that the active zoneset or zone configuration has correctly propagated to the other switches in the fabric.                                                                                                                                                                                             |
| VSAN                      | <b>Interop</b> mode only affects the specified VSAN.                                                                                                                                                                                                                                                                                                                                                                                                  |
| TE ports and PortChannels | TE ports and Port-Channels cannot be used to connect MDS to non-MDS switches. Only E ports can be used to connect to non-MDS switches. TE ports and PortChannels can still be used to connect an MDS to other MDS switches even when in <b>interop</b> mode.                                                                                                                                                                                          |

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

**Table 24-2 Changes in Switch Behavior When Interoperability Is Enabled (continued)**

| Switch Feature                       | Changes if Interoperability Is Enabled                                                                                                                                                             |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FSPF                                 | The routing of frames within the fabric is not changed by the introduction of <b>interop</b> mode. The switch continues to use src-id, dst-id, and ox-id to loadbalance across multiple ISL links. |
| Domain reconfiguration disruptive    | This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing Domain IDs.                                         |
| Domain reconfiguration nondisruptive | This event is limited to the affected VSAN. 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.                                                                                                          |

## Configuring Interoperability

The **interop** mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively. The interoperability procedure is different in Cisco MDS 9500 Series and in 9200 Series switches.



### Note

Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connect 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.

## Cisco MDS 9500 Series Switches

To configure interoperability in a Cisco MDS 9500 Series switch, follow these steps:

- Step 1** Place the VSAN of the E ports (s) that connect to the OEM switch in interoperability mode.

```
switch# config t
switch(config)# vsan database
switch (config-vsan-db)# vsan 1 interop
```

- 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 t
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 does not join the fabric unless the principal switch agrees, and assigns the requested ID.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**



**Note** When changing the Domain ID, the FC IDs assigned to N ports will also change.

**Step 3** Change the Fibre Channel timers (if they have been changed from the system defaults).



**Note** The 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 t
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 MDS domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.

```
switch(config)# fcdomain restart disruptive vsan 1
or
```

- Don't force a fabric reconfiguration

```
switch(config)# fcdomain restart vsan 1
```

## Cisco MDS 9200 Series Switches

To configure interoperability in a Cisco MDS 9200 Series switch, 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
```

**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 t
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 does 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 will also change.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Step 3** Change the Fibre Channel timers (if they have been changed from the system defaults).



**Note**

The 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 t
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 MDS domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.

```
switch(config)# fcdomain restart disruptive vsan 1
```

**or**

- Don't 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.

### Cisco MDS 9500 Series Switches

To verify the resulting status of issuing the interoperability command in a Cisco MDS 9500 Series switch, follow these steps:

**Step 1** Use the **show version** command to verify the version.

```
switch# show ver
Copyright (c) 2001-2005
Cisco Systems, Inc.
Software
 kickstart: version 1.0(3) [gdb]
 System: version 1.0(3) [gdb]
Hardware
 RAM 1932864 kB
 bootflash: 503808 blocks (block size 512b)
 slot0: 0 blocks (block size 512b)
 Compile Time: 2/12/2003 2:00:00
...
```

**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
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

|        |   | 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
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```
switchname MDS9509
username admin password 5 1Li8/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
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```

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 (Seagate) 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 MDS Name Server shows both local and remote entries, and does not timeout the entries.

## Cisco MDS 9200 Series Switches

To verify the resulting status of issuing the interoperability command in a Cisco MDS 9200 Series switch, follow these steps:

**Step 1** Use the **show version** command to verify the version.

```

switch# show ver
Copyright (c) 2001-2005
Cisco Systems, Inc.
Software
 kickstart: version 1.0(3) [gdb]
 System: version 1.0(3) [gdb]
Hardware
 RAM 963116 kB
 bootflash: 503808 blocks (block size 512b)
 slot0: 0 blocks (block size 512b)
 Compile Time: 1/26/2003 2:00:00

```

**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)

```



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

|       |   |      |    |            |    |    |    |
|-------|---|------|----|------------|----|----|----|
| fc1/1 | 1 | auto | on | up         | E  | 2  | -- |
| fc1/2 | 1 | auto | on | fcotAbsent | -- | -- | -- |
| fc1/3 | 1 | auto | on | up         | E  | 2  | -- |
| fc1/4 | 1 | auto | on | down       | -- | -- | -- |
| fc1/5 | 1 | auto | on | down       | -- | -- | -- |
| fc1/6 | 1 | auto | on | up         | E  | 1  | -- |
| fc1/7 | 1 | auto | on | fcotAbsent | -- | -- | -- |
| fc1/8 | 1 | auto | on | fcotAbsent | -- | -- | -- |
| fc1/9 | 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 fc1/1
no shutdown
 interface fc1/2
 interface fc1/3
switchport speed 2000
no shutdown
 interface fc1/4
 interface fc1/5
 interface fc1/6
switchport speed 1000
no shutdown
 interface fc1/7
 interface fc1/8
 interface fc1/9
...
 interface mgmt0
ip address 6.1.1.95 255.255.255.0
no shutdown
vsan database
vsan 1 interop
boot system bootflash:/m9200-system-253e.bin
boot kickstart bootflash:/m9200-kickstart-253e.bin
callhome
fcdomain domain 98 preferred vsan 1
line console
 databits 5
 speed 110
logging linecard
switchname MDS9216
username admin password 5 MF7UQdWLEqUFE 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 state:active
 interoperability mode:yes <----- verify interoperability
 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: 20:01:00:05:30:00:47:9f
 Running fabric name: 10:00:00:60:69:22:32:91
 Running priority: 128
 Current domain ID: 0x62(98) <-----verify domain ID
```

## Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

```

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: 0x62(98) (preferred)
Principal switch run time information:
 Running priority: 2
Interface Role RCF-reject

fc1/1 Upstream Disabled
fc1/3 Non-principal Disabled
fc1/6 Non-principal 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 [Local]
0x63(99) 10:00:00:60:69:c0:0c:1d
0x64(100) 20:01:00:05:30:00:51:1f
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 fc1/1
 1 0x63(99) 500 fc1/3
 1 0x64(100) 1000 fc1/1
 fc1/3
 1 0x65(101) 1000 fc1/6

```

**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 (Seagate) 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

```



## 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, page 25-2](#)
- [Configuring FCS, page 25-3](#)
- [Displaying FCS Information, page 25-4](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## 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 its 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.

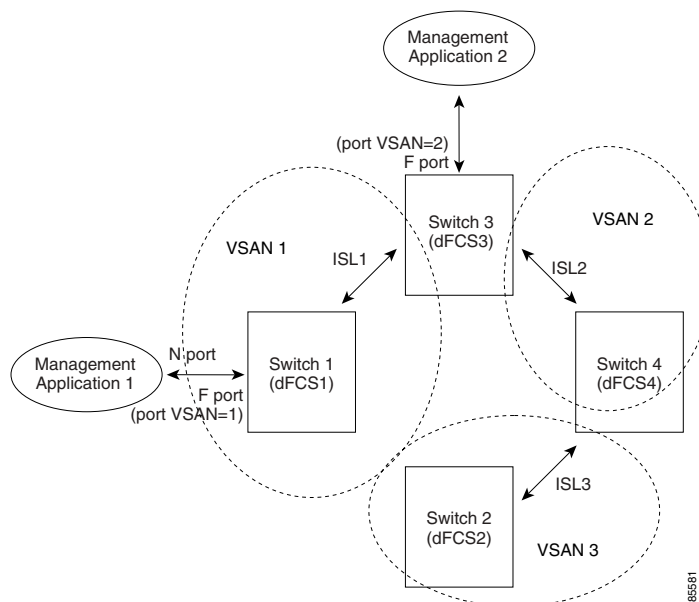
Each object has its own set of attributes and their values. A null value may also be defined for some attributes.

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

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch is a 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 SNMP or CLI.

In Figure 2 Management Application 1 (M1) is connected through an F port with a port VSAN ID 1 and Management Application 2 (M2) is connected through an F port with a port VSAN ID 2. M1 can query FCS info 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 25-1 FCSs in a VSAN Environment**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Significance of FCS

The significance of FCSs are as follows:

- Network Management
  - N port management application can query and obtain information about the fabric elements.
  - SNMP Manager—A SNMP Manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs supports 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 updates it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information, and rebuild its database.
- The SNMP manager can query FCSs for all the IEs, ports, and platforms in the fabric.

## Configuring FCS

Use the **fcs plat-check-global** command to specify if the platform name or node name uniqueness 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# <b>config t</b><br>switch(config)#                                | Enters configuration mode.                           |
| Step 2 | switch(config)# <b>fcs plat-check-global vsan 1</b><br>switch(config)#    | Enables global checking of platform name.            |
|        | switch(config)# <b>no fcs plat-check-global vsan 1</b><br>switch(config)# | Disables (default) global checking of platform name. |

To register platform attributes, follow these steps:

|        | Command                                                                                                                | Purpose                                           |
|--------|------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                                             | Enters configuration mode.                        |
| Step 2 | switch(config)# <b>fcs register</b><br>switch(config-fcs-register)#                                                    | Enters the FCS registration submodule             |
| Step 3 | switch(config-fcs-register)# <b>platform name</b><br><b>SamplePlatform vsan 1</b><br>switch(config-fcs-register-attr)# | Enters the FCS registration attributes submodule. |
|        | switch(config-fcs-register)# <b>no platform name</b><br><b>SamplePlatform vsan 1</b><br>switch(config-fcs-register)#   | Deletes a registered platform.                    |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

|        | Command                                                                                                        | Purpose                                                                                    |
|--------|----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Step 4 | switch(config-fcs-register-attrib)# <b>mgmt-addr 1.1.1.1</b><br>switch(config-fcs-register-attrib)#            | Configures the platform management address.                                                |
|        | switch(config-fcs-register)# <b>no mgmt-addr 1.1.1.1</b><br>switch(config-fcs-register)#                       | Deletes all management addresses on the platform.                                          |
| Step 5 | switch(config-fcs-register-attrib)# <b>nwwn 11:22:33:44:55:66:77:88</b><br>switch(config-fcs-register-attrib)# | Configures the platform node name.                                                         |
|        | switch(config-fcs-register)# <b>no nwwn 11:22:33:44:55:66:77:88</b><br>switch(config-fcs-register)#            | Deletes the platform node name.                                                            |
| Step 6 | switch(config-fcs-register-attrib)# <b>type 5</b><br>switch(config-fcs-register-attrib)#                       | Configures the fc-gs-3 defined platform type.                                              |
|        | switch(config-fcs-register)# <b>no type 5</b><br>switch(config-fcs-register)#                                  | Deletes the configured type and reverts the switch to its factory default of unknown type. |
| Step 7 | switch(config-fcs-register-attrib)# <b>exit</b><br>switch(config-fcs-register)#                                | Exits the FCS registration attributes submode                                              |
| Step 8 | switch(config-fcs-register)# <b>exit</b><br>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 25-1 to 25-9).

### Example 25-1 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
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```
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 25-2 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 25-3 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 25-4 Displays Platform 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 25-5 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 25-6 Displays a List of Switchports in a Specified VSAN

```
switch# show fcs port vsan 24
Port List in VSAN: 24
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```

-- 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 25-7** 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 25-8** 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 25-9** 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

```





## Monitoring System Processes and Logs

---

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying System Processes, page 26-2](#)
- [Displaying System Status, page 26-5](#)
- [Configuring Core and Log Files, page 26-6](#)
- [Configuring HA Policy, page 26-8](#)
- [Resetting HA Statistics, page 26-8](#)
- [Configuring Heartbeat Checks, page 26-8](#)
- [Configuring Watchdog Checks, page 26-8](#)
- [Configuring Upgrade Resets, page 26-9](#)
- [Configuring Kernel Core Dumps, page 26-9](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Displaying System Processes

Use the **show processes** command to obtain general information about all processes (see Examples 26-1 to 26-6).

### Example 26-1 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
```

Terms:

- PID = process ID.
- State = process state
  - D = uninterruptible sleep (usually IO)
  - 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 = how many times a process has been started (or restarted).
- TTY = terminal that controls the process. A “-” usually means a daemon not running on any particular TTY
- Process = name of the process

### Example 26-2 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
```

Terms:

- Runtime(ms) = CPU time the process has used, expressed in milliseconds
- Invoked = number of times the process has been invoked

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

- uSecs = microseconds of CPU time in average for each process invocation
- 1Sec = CPU utilization in percentage for the last one second

### Example 26-3 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
```

Terms:

- 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 26-4 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 FFFFFFFE 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 2AC1F18C4.....*
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

**Example 26-5 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 26-6 Displays Memory Information About Processes**

```
switch# show processes memory
PID MemAlloc StackBase/Ptr Process

1277 120632 7ffffcd0/7ffffefe4 xbar_client
1278 56800 7ffffce0/7fffffb5c wwn
1279 1210220 7ffffce0/7fffffbac vsan
1293 386144 7ffffcf0/7ffffebd4 span
1294 1396892 7ffffce0/7ffffdff4 snmpd
1295 214528 7ffffcf0/7ffff904 rscn
1296 42064 7ffffce0/7fffffb5c qos
```

Where:

- MemAlloc = total memory allocated by the process.
- StackBase/Ptr = process stack base and current stack pointer in hex format

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Displaying System Status

Use the **show system** command to display system-related status information ([Example 26-7](#) to [Example 26-10](#)).

### ***Example 26-7 Displays Default Switch Port States***

```
switch# show system default switchport
System default port state is down
System default trunk mode is on
```

### ***Example 26-8 Displays Error Information for a Specified ID***

```
switch# show system error-id 0x401D0019
Error Facility: module
Error Description: Failed to stop Linecard Async Notifciation.
```

### ***Example 26-9 Displays the System Reset Information***

```
switch# Show system reset-reason
1) No time
 Reason: Watchdog Timeout
 Service:
 Version: 1.0(0.253e)

2) At 125982 usecs after Tue Jan 1 06:45:55 1980
 Reason: Reset Requested CLI command reload
 Service:
 Version: 1.0(0.253e)
```

The **show system reset-reason** command displays the following information:

- In a Cisco MDS 9500 Series 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 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 will not be displayed.

### ***Example 26-10 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 26-11](#)).

### ***Example 26-11 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
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

```
Memory usage: 1027628K total, 313424K used, 714204K free
 3620K buffers, 22278K cache
```

Where:

- Load is defined as 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 shows the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage provides 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.

## Configuring Core and Log Files

You can save cores (from the active supervisor module, the standby supervisor module, or any switching module) to an external flash (slot 0) or to a TFTP server in one of two ways:

- On demand—to copy a single file based on the provided process ID.
- Periodically—to copy core files periodically as configured by the user.

To copy the core and log files on demand, follow this step:

|        | Command                                               | Purpose                                                                              |
|--------|-------------------------------------------------------|--------------------------------------------------------------------------------------|
| Step 1 | switch# <b>copy core:7407 slot0:coreSample</b>        | Copies the core file with the process ID 7407 as coreSample in slot 0.               |
|        | switch# <b>copy core://5/1524 tftp://1.1.1.1/abcd</b> | Copies cores (if any) of a process with pid 1524 generated on slot 5 to tftp server. |

- If the core file for the specified process ID is not available, you will see the following response:

```
switch# copy core:133 slot0:foo
No core file found with pid 133
```

- If two core files exist with same process ID, only one file will be 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# <b>config t</b>                                 | Enters configuration mode.                                                  |
| Step 2 | switch(config)# <b>system cores slot0:coreSample</b>    | Copies the core files coreSample to slot 0.                                 |
|        | switch(config)# <b>system cores tftp://1.1.1.1/abcd</b> | Copies the core files (abcd) in the specified directory on the TFTP server. |
|        | switch(config)# <b>no system cores</b>                  | Disable the core files copying feature.                                     |

A new scheme overwrites any previously-issued scheme. For example, if you issue a new system core command, the cores are periodically saved to the new location or file.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***



**Tip**

Be sure to create any required directory before issuing this command. If the directory specified by this command does not exist, the switch software logs a syslog message each time a copy cores is attempted.)

## 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
```

## Displaying Cores Status

Use the **show system cores** command to display the currently configured scheme for copying cores. See Examples 26-12 to 26-14.

### *Example 26-12 Displays the status of System Cores*

```
switch# show system cores
Transfer of cores is enabled
```

### *Example 26-13 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 Jan 9 03:11
6 fcc 919 Jan 9 03:09
8 acltcam 285 Jan 9 03:09
8 fib 283 Jan 9 03:08
```

Where:

module-num shows the slot number on which the core was generated. In this example, the `fsfp` 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 26-14 Displays Logs on the Local System*

```
switch# show processes log
Process PID Normal-exitStack-traceCore Log-create-time

fsfp 1524 N Y Y Jan 9 03:11
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## Configuring HA Policy

You can disable the HA policy supervisor reset feature (enabled by default) for debugging and troubleshooting purposes.

To configure HA policies, follow this step:

| Step 1 | Command                            | Purpose                                                                                                                             |
|--------|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
|        | switch# <b>system no hap-reset</b> | Disables supervisor reset HA policy.                                                                                                |
|        | switch# <b>system hap-reset</b>    | Enables Supervisor Reset HA policy whenever a critical service runs out of HA policies (default) and reverts it to factory default. |

## Resetting HA Statistics

The system statistics reset feature resets the high availability statistics collected by the system.

```
switch# system statistics reset
```

## Configuring Heartbeat Checks

The software monitors every service to verify if heartbeats are sent at regular intervals. If not, the software restarts that service. This feature helps locate situations when a service is stuck in an infinite loop.

You can disable the heartbeat checking feature (enabled by default) for debugging and troubleshooting purposes like attaching a GDB to a specified process.

To configure heartbeat checks, follow this step:

| Step 1 | Command                            | Purpose                                                               |
|--------|------------------------------------|-----------------------------------------------------------------------|
|        | switch# <b>system no heartbeat</b> | Disables heartbeat checks.                                            |
|        | switch# <b>system heartbeat</b>    | Enables heartbeat checks (default) and reverts it to factory default. |

## Configuring Watchdog Checks

If a watchdog is not logged at every 8 seconds by the software, the supervisor module reboots the switch.

You can disable the watchdog checking feature (enabled by default) for debugging and troubleshooting purposes like attaching a GDB or a kernel GDB (KGDB) to a specified process.

To configure watchdog checks, follow this step:

| Step 1 | Command                           | Purpose                                                              |
|--------|-----------------------------------|----------------------------------------------------------------------|
|        | switch# <b>system no watchdog</b> | Disables watchdog checks.                                            |
|        | switch# <b>system watchdog</b>    | Enables watchdog checks (default) and reverts it to factory default. |



*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Configuring Upgrade Resets

This feature enables supervisor module resets when an upgrade has failed. If the upgrade fails for any reason, the software reboots the switch since the file system may be in an unstable state.

You can disable the upgrade-reset feature (enabled by default) for debugging and troubleshooting purposes.

To configure supervisor upgrade resets, follow this step:

|        | Command                                | Purpose                                                                        |
|--------|----------------------------------------|--------------------------------------------------------------------------------|
| Step 1 | switch# <b>system no upgrade-reset</b> | Disables the upgrade reset feature.                                            |
|        | switch# <b>system upgrade-reset</b>    | Enables the upgrade reset feature (default) and reverts it to factory default. |

## Configuring 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 1GB 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.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

To configure the external server, follow these steps:

|        | Command                                                          | Purpose                                      |
|--------|------------------------------------------------------------------|----------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                | Enters configuration mode.                   |
| Step 2 | switch(config)# <b>kernel core target 10.50.5.5</b><br>succeeded | Configures the external server's IP address. |

To configure the module information, follow these steps:

|        | Command                                                               | Purpose                                                                                          |
|--------|-----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                     | Enters configuration mode.                                                                       |
| Step 2 | switch(config)# <b>kernel core module 5</b><br>succeeded              | Configures kernel core generation for module 5.                                                  |
|        | switch(config)# <b>kernel core module 5 level header</b><br>succeeded | Configures kernel core generation for module 5, and limits the generation to header-level cores. |
| Step 3 | switch(config)# <b>kernel core limit 2</b><br>succeeded               | Configures generations for two modules. The default is 1 module.                                 |

All changes made to kernel cores are saved to the running configuration and may be viewed using the **show running-config** command. Alternatively, use the **show kernel cores** command to view specific configuration changes (see examples 26-15 to 26-17).

#### ***Example 26-15 Displays the Core Limit***

```
switch# show kernel core limit
2
```

#### ***Example 26-16 Displays the External Server***

```
switch# show kernel core target
10.50.5.5
```

#### ***Example 26-17 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
```



---

## Numerics

### 16-port modules

BB\_credits [9-10](#)

LEDs [9-12](#)

preserving configurations [6-6](#)

See also switching modules

### 32-port modules

BB\_credits [9-10](#)

configuration guidelines [9-7](#)

preserving configurations [6-6](#)

SPAN guidelines [23-6](#)

See also switching modules

---

## A

### AAA

authorization and authentication process [14-4](#)

setting authentication [14-5](#)

usage [1-8](#)

user accountability [14-11](#)

### access control

iSCSI enforcement [17-50](#)

iSCSI mechanisms [17-49](#)

accounting [14-11](#)

### active

modules [6-2](#)

states [4-7](#)

zones [12-6](#)

zone sets [12-11](#)

### adding

IP addresses [16-14](#)

ports to a PortChannel [11-6](#)

SNMP communities [14-24](#)

switches [3-4](#)

address-allocation cache [19-14](#)

Address Resolution Protocol

See ARP

administrative speed

configuring [9-9](#)

administrative states

description [9-5](#)

administrator passwords

configuring [3-9](#)

configuring switch [3-3](#)

creating additional accounts [3-5](#)

default [3-5](#)

recovering [14-13](#)

requirements (note) [3-6, 3-10](#)

advertisement packets

setting time intervals [16-15](#)

aggregated flow statistics [15-12](#)

aliases

configuring [12-4](#)

ALPA [9-22](#)

ARP

clearing and viewing entries [16-6](#)

IP services [1-5](#)

ARP caches

clearing [17-8](#)

managing [17-8](#)

assigning

alias names [12-4](#)

contact information [18-2](#)

domain IDs [19-4](#)

FC IDs [15-8](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- global keys [14-15](#)
  - host key [14-14](#)
  - users [14-24](#)
  - zone members [12-4](#)
  - authentication
    - CHAP option [17-65](#)
    - global override [17-51](#)
    - iSCSI hosts [17-51](#)
    - iSCSI setup [17-64](#)
    - See MD5 authentication
    - See simple text authentication
  - authentication, authorization, and accounting
    - See AAA
  - automatic synchronization
    - conditions [4-7](#)
    - modules [4-6](#)
  - AutoNotify
    - destination profile (note) [18-5](#)
    - registration [18-3](#)
    - service contract [18-2](#)
  - auto port mode
    - configuring [9-9](#)
    - description [9-5](#)
    - interface configuration [9-2](#)
- 
- B**
- basic input/output system
    - See BIOS
  - BB\_credits
    - configuring [9-10](#)
    - reason codes [9-6](#)
  - beacon mode
    - configuring [9-11](#)
    - identifying LEDs [9-12](#)
    - LEDs [6-10](#)
  - Berkeley Packet Filter
    - See BPF
  - BIOS
    - boot sequence [5-25](#)
    - recovering corrupted bootflash [5-27](#)
    - recovery sequence [5-26](#)
    - setup (figure) [5-28](#)
  - BIOS upgrades [5-16](#)
  - boot
    - sequence [5-25](#)
    - variable synchronization [4-6](#)
  - bootflash
    - copying to [5-12, 5-21](#)
    - description [2-12](#)
    - device [5-12](#)
    - file system [5-1](#)
    - initializing [3-23](#)
    - recovering corrupted [5-25 to 5-26](#)
    - space requirements [5-2](#)
    - See also internal bootflash
  - bootloader
    - loading kickstart [5-25](#)
    - nondisruptive upgrades [5-15](#)
    - skipping phases [5-30](#)
  - bootup diagnostics [6-3](#)
  - boot variable
    - upgrading single supervisor switches [5-21](#)
  - boot variables
    - disruptive upgrades [5-24](#)
    - specifying [5-18](#)
  - BPF
    - library [24-14](#)
    - See also libpcap freeware
  - B ports
    - functionality [9-4](#)
  - broadcast
    - in-band addresses default [6-10](#)
    - routing [15-9](#)
  - buffer-to-buffer credits
    - See BB\_credits
  - build fabric frames [19-3](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## C

### cache

See address-allocation cache

### Call Home

configuring [18-3 to 18-7](#)

### Call Home

functionality [1-6](#)

message format options [18-2](#)

### capture filters [24-14](#)

### CDP

clearing [3-34](#)

configuring [3-33](#)

configuring globally [3-34](#)

displaying [3-35](#)

hold time [3-34](#)

packet transmission [3-33](#)

### CHAP authentication [17-65](#)

### chassis

types [6-1](#)

### checks

See compatibility checks

See heartbeat checks

See watchdog checks

### Cisco Discovery Protocol

See CDP

### Cisco MDS 9000 System Debug Server [26-9](#)

### Cisco MDS 9200 Series

configuring interoperability [24-22](#)

LEDs [6-8](#)

mgmt0 LEDs [6-9](#)

supervisor modules [6-2](#)

verifying interoperability [24-26](#)

### Cisco MDS 9216 switches

high availability [1-3, 4-2](#)

modules [1-7, 6-1](#)

overview [1-1](#)

supervisor module [6-4](#)

### Cisco MDS 9500 Series

configuring interoperability [24-21](#)

high availability [1-3, 4-2](#)

LEDs [6-9](#)

overview [1-1](#)

supervisor modules [6-2](#)

verifying interoperability [24-23](#)

### Cisco MDS 9506 Directors

modules [1-7, 6-1](#)

overview [1-1, 1-2](#)

### Cisco MDS 9509 Directors

modules [1-7, 6-1](#)

overview [1-1, 1-2](#)

### clearing

FIB statistics [15-13](#)

FSPF counters [15-8](#)

zone sets [12-11](#)

### CLI

accessing submodes [2-3](#)

alternative [1-9](#)

command modes [2-3](#)

syslog provisioning [1-8](#)

updating SNMPv3 passwords [14-24](#)

### clock modules

monitoring status [7-10](#)

### CMOS

configuration [5-28](#)

saving changes [5-29](#)

### COM 1

configuring [3-32](#)

### command-line interface

See CLI

### commands

saving output to files [2-18](#)

### CompactFlash

devices [3-22, 3-23, 5-12](#)

disk [5-1](#)

slot 0 [5-12](#)

### compatibility checks [11-7](#)

### computing routes [15-1](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## configuring

FCIP links [17-20](#)

## congestion control methods

See edge quench congestion control

See FCC

## connecting a modem

COM 1 [3-32](#)

## consistent switch states [11-6](#)

## console port

logging in [5-20](#)

## console session

severity levels [21-5](#)

## control frames [17-18](#)

## control traffic

disabling [20-4](#)

## core dumps

full [17-16](#)

IPS [17-16](#)

kernel [17-16, 26-9](#)

partial [17-16](#)

## cores [26-6](#)

## customized

targets [22-4](#)

## D

## databases

See zone databases

## data field

configuring size [9-11](#)

## data frames [17-18](#)

## dead time interval [15-6](#)

## default gateway

BIOS setup configuration [5-28](#)

configuring mgmt0 Ethernet interfaces [9-14](#)

recovering loader> prompt [5-30](#)

recovering switch(boot)# prompt [5-31](#)

## default groups [14-23](#)

## default zones

description [12-9](#)

interoperability [24-20](#)

## deleting

FSPF configurations [15-4](#)

PortChannels [11-6](#)

## destination IDs

exchange based [11-5](#)

flow based [11-4](#)

frame identification [20-2](#)

frame loop back [24-3](#)

in-order delivery [15-9, 20-2](#)

load balancing [1-5, 11-1](#)

path selection [8-7](#)

## destination profiles

configuring [18-5](#)

## device IDs

Call Home format [18-12, 18-13](#)

copying files [14-23](#)

report capacity [22-1](#)

## Device View

description [1-9](#)

## digital signature algorithm

See DSA key pairs

## Dijkstra's algorithm [15-2](#)

## direct memory access

See DMA-bridge

## disabling routing protocols [15-4](#)

## discovered

LUNs [22-3](#)

targets [22-2](#)

## display filters

selective viewing [24-10](#)

## disruptive

switchover [4-4](#)

upgrades [5-4](#)

## distribution tree [15-9](#)

## DMA-bridge [17-9](#)

## documentation

related documents [xxiv](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## domain IDs

- configuring [19-4](#)
- distributing [19-2](#)
- failure [9-7](#)
- interoperability [24-20](#)
- preferred [19-5](#)
- range [2-20](#)
- static [19-5](#)

## domain manager

- isolation [9-7](#)

## domain names

- defining [16-19](#)

## Domain Name System servers

- See DNS servers

## domain overlap

- isolation [9-7](#)

## drop latency time

- configuring [15-11](#)

## dsa key pairs

- generating [14-19](#)

## E

### edge quench congestion control

- description [20-2](#)

### egress port [23-12](#)

### EISL

- functionality [1-4](#)
- PortChannel links [11-1](#)

### ELP failure [9-7](#)

### e-mail notification

- Call Home [18-1](#)

### environmental monitors [6-8, 6-9](#)

### E ports

- 32-port guidelines [9-7](#)
- classes of service [9-3](#)
- configuring [9-9](#)
- FSPF topology [15-2](#)
- interface modes [9-2](#)

### isolation [9-7](#)

### recovering from isolation [12-10](#)

### SPAN [23-3](#)

### trunking [1-4](#)

### trunking configuration [10-3](#)

### error disabled code [9-6](#)

### error messages

#### description [21-2](#)

### error state [5-32](#)

### ESC failure [9-7](#)

### Ethereal freeware

- analyzer [24-6](#)
- information [24-5](#)

### Ethernet PortChannels

#### configuring [17-14](#)

### exchange IDs

- in-order delivery [15-9](#)
- load balancing [1-5, 11-1, 24-3](#)
- path selection [8-7](#)

### exchange link parameter

- See ELP failure

### exporting

- zone databases [12-10](#)

### extended ISL

- See EISL

### external RADIUS server

- CHAP [17-65](#)

### external server

- configuring [26-10](#)

## F

### fabric

- See build fabric frames
- See reconfigure fabric frames

### Fabric Analyzer

- capture range [2-20](#)
- configuring [24-7](#)
- description [24-5](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- frame range [2-20](#)
- Fabric Configuration Server
  - See FCS
- fabric login
  - See FLOGI
- Fabric Manager
  - description [1-8](#)
  - Device View [1-9](#)
  - Fabric View [1-9](#)
- fabric names
  - setting [19-8](#)
- fabric pWWNs
  - configuring zones [12-4](#)
  - zone membership [12-2](#)
- fabric reconfiguration
  - fcdomain phase [19-2](#)
- fabric shortest path first
  - See FSPF
- Fabric View
  - description [1-9](#)
- fail over protection [17-13](#)
- fan modules
  - monitoring status [7-10](#)
- fan trays
  - overview [1-2](#)
- fault tolerant fabric
  - example (figure) [15-2](#)
- FC aliases
  - configuring zones [12-4](#)
- fcanalyzer
  - clearing hosts [24-9](#)
  - displaying filters [24-10](#)
- FCC
  - benefits [20-2](#)
  - default settings [20-4](#)
  - enabling [20-3](#)
  - frame handling [20-2](#)
  - logging facility [21-2](#)
- fcdomain
  - configuring [19-1](#)
  - default settings [19-14](#)
- FC IDs
  - address format [2-20](#)
  - allocating [19-2, 24-18](#)
  - allocating areas [24-18](#)
  - configuring zones [12-4](#)
- FCIP
  - configuring [17-17](#)
  - Gigabit Ethernet ports [17-4](#)
  - interfaces [17-18, 17-19](#)
  - IPS module [17-2, 17-17](#)
  - IP storage [17-1](#)
  - links [17-18](#)
  - profiles [17-18, 17-19](#)
  - TCP connections [17-18](#)
- FCIP parameters
  - default [17-76](#)
- Fcot not present [9-6](#)
- fcping
  - invoking [24-4](#)
- FCS
  - configuring [25-3](#)
  - description [25-2](#)
  - logging facility [21-2](#)
  - significance [25-3](#)
- fctrace
  - invoking [24-3](#)
- Fibre Channel analyzers [23-10](#)
- Fibre Channel Congestion Control
  - See FCC
- Fibre Channel domain
  - See fcdomain
- Fibre Channel over IP
  - See FCIP
- Fibre Channel traffic
  - SPAN sources [23-3](#)
- file system
  - formatting [3-23](#)



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- redirection [2-18](#)
- volatile [2-12](#)
- File Transfer Protocol
  - See FTP
- filters
  - capture [24-14](#)
  - defining display [24-11](#)
- FLOGI
  - displaying details [13-1](#)
  - logging facility [21-2](#)
- flow statistics [15-12](#)
- FL ports
  - classes of service [9-3](#)
  - configuring [9-9](#)
  - fctrace [24-3](#)
  - interface modes [9-2](#)
  - nonparticipating code [9-7](#)
  - persistent FC IDs [19-10](#)
  - SPAN [23-3](#)
- F ports
  - classes of service [9-3](#)
  - configuring [9-9](#)
  - interface modes [9-2](#)
  - SPAN [23-3](#)
- frames
  - control [17-18](#)
  - data [17-18](#)
  - encapsulation [23-8](#)
  - flow [1-7](#)
  - MTU [17-5](#)
  - reordering [15-9](#)
- FSPF
  - alternative paths [15-1](#)
  - clearing counters [15-8](#)
  - computing link cost [15-5](#)
  - configuring globally [15-3](#)
  - configuring on interfaces [15-5](#)
  - default settings [15-18](#)
  - disabling on interfaces [15-6](#)
  - disabling routing protocols [15-4](#)
  - hold time range [2-20, 15-1](#)
  - interoperability [24-21](#)
  - link state protocol [15-2](#)
  - reconvergence time [15-2](#)
  - routing services [15-1](#)
  - topologies example [15-2](#)
- FTP
  - logging facility [21-2, 21-8](#)
- full zone set
  - considerations [12-6](#)
  - distribution [12-11](#)
- Fx ports
  - 32-port default [9-8](#)
  - configuring [9-9](#)
  - FCS [25-2](#)
  - interface modes [9-4](#)

---

## G

- Gigabit Ethernet
  - configuring [17-5](#)
  - IP routing [17-7](#)
  - major interfaces [17-6](#)
  - ports [17-4](#)
  - subinterfaces [17-6](#)
- Gigabit Ethernet parameters
  - default [17-76](#)
- global iSCSI information
  - displaying [17-55](#)

---

## H

- HA policy [26-8](#)
- hardware
  - displaying inventory [7-2](#)
  - status description [6-3](#)
- hard zoning [12-5](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

HA-standby [4-5, 6-2](#)  
 HA switchover [4-3, 4-6](#)  
 heartbeat checks [26-8](#)  
 hello time interval [15-5](#)  
 hidden routes [15-14](#)  
 high availability  
   default setting [4-8](#)  
   Ethernet PortChannel [17-64](#)  
   features [17-62](#)  
   functionality [1-3, 4-2](#)  
   process restartability [4-5](#)  
   software upgrade [5-4](#)  
   status [4-6](#)  
   VRRP [17-63](#)  
   VRRP features [17-13](#)  
   See also HA policy  
   See also HA standby  
   See also HA switchover

---

ICMP statistics  
   displaying [17-11](#)  
 identical passwords  
   CLI and SNMP [14-24](#)  
 IDs  
   CCO IDs [18-3](#)  
   contract IDs [18-4, 18-12](#)  
   customer IDs [18-4](#)  
   image version and IDs [5-1](#)  
   login IDs [3-5](#)  
   process IDs [3-28, 26-2, 26-6](#)  
   profile IDs [18-5](#)  
   region ID [15-4](#)  
   serial IDs [18-13](#)  
   server IDs [18-13](#)  
   site IDs [18-4, 18-12](#)  
   See destination IDs  
   See device IDs

See domain IDs  
 See exchange IDs  
 See FC IDs  
 See port IDs  
 See source IDs  
 See user IDs  
 See VR IDs  
 See VSAN IDs  
 images  
   See kickstart images  
   See software images  
   See system images  
 software upgrades  
   See also one-step upgrade  
   See also step-by-step upgrades  
 importing database [12-10](#)  
 inactive code [9-6](#)  
 inconsistent switch states [11-6](#)  
 ingress port [23-10](#)  
 initiator access list [17-49](#)  
 in-order delivery [15-9](#)  
   enabling [15-11](#)  
 in-order guarantee [15-10](#)  
 install all  
   command benefits [5-5](#)  
   command examples [5-8](#)  
   command failure cases [5-6](#)  
   command function [5-5](#)  
   command requirements [5-3](#)  
   command usage [5-7](#)  
   remote location path (caution) [5-9](#)  
   TFTP get operation [5-10](#)  
 insufficient power [6-3](#)  
 interface  
   adding to PortChannels [11-7](#)  
   configuring FSPF [15-5](#)  
   suspended states [11-7](#)  
 interfaces  
   characteristics [9-2](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- configuring [9-8](#)
- data field size [9-11](#)
- default settings [9-13](#)
- description [9-10](#)
- modes [9-2, 9-9](#)
- reason codes [9-5](#)
- states [9-5](#)
- internal bootflash
  - description [2-12](#)
  - flash devices [2-12, 3-22](#)
  - See also bootflash
- internal switch states
  - description [4-8](#)
- interoperability
  - configuring [24-20](#)
  - verifying status [24-23 to 24-28](#)
- inter-switch links
  - See ISL
- invoking fcping [24-4](#)
- IP address
  - address format [2-20](#)
  - SMTP server [18-7](#)
- IP addresses
  - configuring in VSANs [16-5](#)
- IPFC
  - logging facility [21-3](#)
- IP features
  - default settings [16-20](#)
- IP forwarding
  - disabling [16-5](#)
- IP over Fibre Channel
  - See IPFC
- IP routing
  - Gigabit Ethernet interface [17-7](#)
  - static [1-5](#)
- IPS core dumps
  - See core dumps
- IP services
  - default settings [16-19](#)
- IPS module
  - CDP [3-33](#)
  - CDP support [17-16](#)
  - functions [17-2](#)
  - port mode [17-4](#)
- IPS Ports
  - multiple connections [17-62](#)
- IPS ports [17-2](#)
- IPS services module
  - See IPS module
- IPS statistics
  - displaying [17-60](#)
- IP statistics
  - displaying [17-10](#)
- IP storage
  - VRRP [17-13](#)
- iSCSI
  - Gigabit Ethernet ports [17-4](#)
  - IPS module [17-2](#)
  - node name [17-46](#)
- iSCSI authentication [17-51](#)
- iSCSI discovery [17-50](#)
- iSCSI host
  - multiple VSANs [17-48](#)
- iSCSI hosts
  - configuring accessibility [17-49](#)
- iSCSI initiators
  - assigning WWNs [17-48](#)
  - displaying [17-56](#)
  - dynamic mapping [17-46](#)
  - identifying [17-46](#)
- iSCSI interfaces
  - displaying [17-53](#)
- iSCSI parameters
  - default [17-77](#)
- iSCSI session creation [17-50](#)
- iSCSI sessions
  - displaying [17-55](#)
- iSCSI targets

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

access control [17-49](#)  
 secondary access [17-43](#)  
 iSCSI user information  
   displaying [17-61](#)  
 iSCSI virtual targets  
   displaying [17-59](#)  
 ISL  
   PortChannel links [11-1](#)  
 isolation  
   reason codes [9-7](#)

---

## J

jumbo frames  
   see MTU frame size

---

## K

kernel core dumps [26-9](#)  
   configuring [26-10](#)  
 kickstart images  
   downloading [3-23](#)  
   KICKSTART variable [5-2](#)  
   loading system images [5-25](#)  
   overview [5-1](#)  
   recovering corrupted [5-30](#)  
   recovery [5-31](#)  
   recovery interruption [5-26](#)  
   specifying [5-18](#)  
   verifying integrity [5-14](#)

---

## L

LEDs  
   identifying beacon [9-12](#)  
 libpcap freeware [24-5](#)  
 link cost [15-2](#)  
 link end points [17-18](#)

link failure [9-6](#)  
   high availability [4-2](#)  
 link redundancy  
   Ethernet PortChannels [17-14](#)  
 load balancing [11-1](#)  
   attributes [8-7](#)  
   guarantee [8-8](#)  
   mechanisms [11-4](#)  
 local capture [24-7](#)  
 log files [26-6](#)  
   configuring [21-6](#)  
 logging  
   default settings [21-13](#)  
   severity levels [21-4](#)  
   system messages [21-2](#)  
 logical unit numbers  
   See LUNs  
 loop monitoring [24-18](#)  
 loop port [24-18](#)  
 LSR [15-16](#)  
 LUNs  
   address format [2-20](#)  
   displaying discovered, example [22-3](#)

---

## M

MAC address  
   format [2-20](#)  
 major threshold [6-8](#)  
 Management Information Base  
   See MIB  
 management module [6-8](#)  
 management redundancy  
   high availability [4-2](#)  
 manual assignment [17-47](#)  
 mapping  
   iSCSI hosts [17-46](#)  
 MD5 authentication [16-16](#)  
 memory test [5-26, 5-27](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

mgmt0 interfaces

- autosensing port [9-14](#)
- configuring [9-14](#)
- configuring ethernet ports [16-3](#)
- overview [9-1](#)
- recovery from switch(boot)# prompt [5-31](#)

minor threshold [6-8](#)

modify existing users [14-23](#)

module

- configuring logging [21-6](#)

module configuration

- sample scenarios [6-6](#)

module status [9-1](#)

module temperature [7-9](#)

monitoring traffic [23-6](#)

MTU frame size [17-5](#)

multicast routing [15-9](#)

multi-pid option [13-7](#)

## N

name server

- interoperability [24-21](#)

name server proxy [13-3](#)

network administrator [2-19](#)

network operator [2-19](#)

Network Time Protocol

- See NTP

network traffic

- monitoring [23-6](#)

next hop domain ID [15-8](#)

NL ports

- fctrace [24-3](#)
- interface modes [9-5](#)
- zone enforcement [12-5](#)

node WWNs

- See nWWNs

nondisruptive

- restart [4-2](#)

- switchover [4-4](#)
- upgrades [5-4](#)

nonparticipating code [9-7](#)

non-trunking ISL [10-2](#)

nonvolatile storage [6-6](#)

N ports

- fctrace [24-3](#)
- zone enforcement [12-5](#)
- zone membership [12-2](#)

NTP

- logging facility [21-3](#)

nWWNs

- address format [2-20](#)

Nx ports

- hard zoning [12-5](#)

## O

offline code [9-6](#)

one-step upgrade

- install all command [5-4](#)
- reload command [5-4](#)

operational interfaces

- viewing PortChannels [11-9](#)

operational state [9-9](#)

operational state setting

- description [9-5](#)

originator exchange IDs

- See exchange IDs

out-of-order delivery [15-9](#)

## P

password recovery [14-13](#)

path discovery [24-3](#)

permitted filters [24-15](#)

persistent FC ID [19-9](#)

persistent FC IDs

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- displaying [19-12](#)
- physical interfaces [17-10](#)
- port aggregation [4-2](#)
- PortChannel
  - configuring FC routes [15-8](#)
  - functionality [1-5](#)
  - high availability [4-2](#)
  - in-order guarantee [15-10](#)
  - link changes [15-10](#)
  - link failure [15-3](#)
  - load balancing [1-5](#)
  - logging facility [21-3](#)
  - membership [11-8](#)
  - range [2-20](#)
  - reason codes [9-7](#)
- PortChannels
  - adding Gigabit Ethernet interfaces [17-16](#)
  - adding interfaces [11-6](#)
  - configuring [11-5](#)
  - default settings [11-11](#)
  - deleting [11-6](#)
  - examples [11-2](#)
  - forcing additions [11-7](#)
  - guidelines [11-8](#)
  - interoperability [24-20](#)
  - member combinations [17-15](#)
  - SPAN [23-3](#)
  - trunking comparison [11-3](#)
- port group [9-7](#)
- port IDs
  - configuring zones [12-4](#)
- port mode
  - IPS [17-2, 17-4](#)
- port modes
  - auto [9-5](#)
- ports
  - virtual E [17-18](#)
- Port world wide name
  - See pWWN
- port WWNs
  - See pWWNs
- power supplies [1-2, 6-8, 6-9](#)
  - configuring [7-6](#)
  - displaying configuration [7-6](#)
  - guidelines [7-6](#)
  - modes [3-30, 7-6](#)
- power usage
  - displaying details [7-5](#)
- preempt option [16-16](#)
- preferred domain IDs [19-5](#)
- preshared key [14-15](#)
- principle switch [19-4, 19-5](#)
  - selecting [19-1](#)
- private device [9-22](#)
- process ID [26-6](#)
- Process Logs [26-4](#)
- process restartability [4-5](#)
- protocol analysis [24-5](#)
- pWWNs
  - address format [2-20](#)
  - configuring zones [12-4](#)
  - zone membership [12-2](#)

---

## Q

### QoS

- default settings [20-4](#)
- displaying information [20-4](#)
- enabling control traffic [20-4](#)
- logging facilities [21-3](#)
- priority queuing [1-7](#)
- quality of service
  - See QoS

---

## R

- R\_A\_TOV time [9-6](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## RADIUS

- AAA solutions [1-8](#)
- authorization process [14-17](#)
- configured parameters [14-18](#)
- secret key [1-8](#)
- setting preshared key [14-15](#)
- specifying servers [14-14](#)
- specifying time-out [14-15](#)
- rebooting switch [6-5](#)
- rebooting system [5-19](#)
- reconfigure fabric [9-6](#)
- reconfigure fabric frames [19-3](#)
- reconvergence time
  - FSPF [15-2](#)
- recovering passwords [14-13](#)
- recovery sequence [5-26](#)
- redundancy states [4-7](#)
- redundant physical links [15-3](#)
- Registered State Change Notification
  - See RSCN
- remote capture [24-7, 24-9](#)
- remote capture daemon [24-6](#)
- Remote Capture Protocol
  - See RPCAP
- Remote Monitoring
  - See RMON
- retransmit intervals [15-7](#)
- route cost
  - computing [15-5](#)
- route table [17-7](#)
- routing
  - See broadcast routing
  - See IP routing
- RPCAP
  - Ethereal communication [24-6](#)
- rsa1 key pairs
  - generating [14-19](#)
- rsa key pairs
  - generating [14-19](#)

## RSCN

- logging facility [21-3](#)
- run time checks [15-7](#)

---

## S

- SAN operating system
  - See SAN-OS
- SAN-OS [5-2](#)
- SCSI LUNs
  - discovering targets [22-1](#)
- SD ports
  - bidirectional traffic [23-12](#)
  - configuring [9-9, 23-6](#)
  - interface modes [9-2, 9-4](#)
- secondary MAC address [24-16](#)
- Secure Shell
  - See SSH
- security features
  - default settings [14-26](#)
- security parameter index
  - See SPI
- See MAC address
  - See also WWNs
- selective purging
  - persistent FC IDs [19-11](#)
- severity levels
  - logging [21-6](#)
- shutdown state [9-7](#)
- Simple Network Management Protocol
  - See SNMP
- simple text authentication [16-16](#)
- simulating
  - Call Home [18-8](#)
- slot0
  - formatting [3-23](#)
- small computer system interface
  - See SCSI
- SMARTnet [18-3](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

## SMTP

server address [18-6](#)

## SNMP

access control [14-21](#)

access groups [14-21](#)

CLI configuration [14-23](#)

community strings [14-21](#)

configuring from CLI [14-23](#)

counter Information [14-25](#)

displaying information [14-25](#)

groups [14-22](#)

read-write access [14-24](#)

server contact [18-2](#)

versions [14-20](#)

## SNMP manager

FCS [25-3](#)

## SNMPv3

security features [14-21](#)

## software image

default setting [5-34](#)

error state [5-25](#)

recognizing errors [5-33](#)

startup configuration [5-19](#)

## software images

bootflash corruption [5-25](#)

compatibility issues [5-15](#)

corruption [5-25](#)

recovery procedure [5-26](#)

saving [5-19](#)

space requirement [5-2](#)

specifying [5-18](#)

synchronizing [4-6](#)

upgrade requirements [5-2](#)

upgrading [5-1](#)

variables [5-2](#)

## software upgrades

disruptive [5-7](#)

high availability [4-2](#)

manual, dual supervisor [5-12](#)

mechanisms [5-4](#)

quick [5-24](#)

soft zoning [12-5](#)

## source IDs

Call Home event format [18-12](#)

exchange based [11-5](#)

flow based [11-4](#)

frame identification [20-2](#)

frame loop back [24-3](#)

in-order delivery [15-9](#)

load balancing [1-5, 11-1](#)

path selection [8-7](#)

## SPAN

configuring sessions [23-5](#)

default settings [23-14](#)

egress source [23-3](#)

encapsulating frames [23-8](#)

FC analyzers [23-10](#)

ingress source [23-2](#)

monitoring traffic [23-2](#)

source configuration [23-4](#)

sources [23-3](#)

## speed

LEDs [6-10](#)

## SPI

configuring virtual router [16-16](#)

## SSH

default service [14-18](#)

force option [14-19](#)

host key pair [14-19](#)

protocol status [14-20](#)

session [5-12, 5-20](#)

## SSH session

message logging [21-5](#)

standby module [6-2](#)

monitoring [4-2](#)

standby supervisor [4-6](#)

## stateful

HA-switchover [4-3](#)



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- stateless
  - warm switchover [4-3](#)
- static domain IDs [19-5](#)
- static mapping
  - iSCSI initiators [17-47](#)
  - pWWN assignment [17-47](#)
- static routes [8-9](#)
  - run time checks [15-7](#)
- status
  - LEDs [6-10](#)
- step-by-step upgrade [5-4](#)
- storage
  - permanent and temporary [2-12](#)
- storage devices
  - access control [12-1](#)
- subnet mask
  - BIOS setup configuration [5-28](#)
  - configuring IP routes [16-6](#)
  - configuring mgmt0 [3-20](#)
  - configuring mgmt0 interfaces [9-14, 16-2](#)
  - configuring switch [3-3](#)
  - default setting [6-10](#)
  - initial configuration [3-6, 3-10](#)
  - loader> prompt recovery [5-30](#)
  - switch(boot)# prompt recovery [5-31](#)
- subnetwork requirements [17-6](#)
- subordinate switch [19-7](#)
- supervisor module
  - CDP support [17-16](#)
  - default settings [6-10](#)
- supervisor modules
  - active [1-7, 4-2](#)
  - active state [4-7, 4-8, 6-3](#)
  - automatic synchronization [4-5](#)
  - default settings [6-10](#)
  - dual modules [6-2](#)
  - high availability [4-2](#)
  - major threshold [6-9](#)
  - recovering password [5-33](#)
  - resetting [6-5](#)
  - standby module [1-7](#)
  - standby state [4-7, 4-8, 6-3](#)
  - standby status [4-5, 6-3](#)
  - states [4-7](#)
  - switch options [1-7](#)
  - switchover [4-3](#)
  - synchronizing images [4-5](#)
  - upgrading [5-20](#)
  - viewing information [6-4](#)
- suspended state [11-7](#)
- switch
  - dual supervisor [5-33](#)
  - reliability service [1-3](#)
  - reloading [6-5](#)
  - role-based access [1-8](#)
  - secure access [1-8](#)
  - security management [1-8](#)
  - single supervisor [5-31](#)
  - SNMPv3 access [1-8](#)
  - verifying modules [6-2](#)
- switchability
  - high availability [4-2](#)
- switched port analyzer
  - See SPAN
- switching module
  - 16-port [6-6](#)
  - 32-port [6-6](#)
  - image [6-2](#)
  - LEDs [6-8](#)
  - power cycle [6-5](#)
  - powering off [6-7](#)
  - reloading [6-5](#)
  - status [6-2](#)
  - viewing states [6-3](#)
- switching modules
  - connecting to [5-17, 6-4](#)
  - LEDs [6-8](#)
  - LEDs (table) [6-10](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- managing [6-1](#)
- powering off [6-7](#)
- preserving configuration [6-6](#)
- progression states [6-3](#)
- reloading [6-5](#)
- reset [4-3](#)
- resetting [6-5](#)
- states [6-1](#)
- thresholds [7-9](#)
- switchover mechanism
  - HA [4-5, 4-7, 4-8, 6-3](#)
  - warm [4-5, 4-7, 4-8, 6-3](#)
- switch priority
  - configuring [19-6](#)
  - range [2-20](#)
- switch redundancy states [4-7](#)
- switch states [11-6](#)
- synchronization
  - See automatic synchronization
- syslogs
  - viewing [1-8](#)
- syslog server [21-2](#)
  - configuring [21-7](#)
- system assignment [17-47](#)
- system image [3-23](#)
  - reading configuration [5-25](#)
  - recovery interruption [5-26](#)
  - specifying [5-18](#)
  - switching module [6-2](#)
- system images [5-1](#)
  - specifying [5-18](#)
  - SYSTEM variable [5-2](#)
- system messages
  - configuring [21-5](#)
  - default settings [21-13](#)
  - displaying configuration [21-8](#)
  - format [21-4](#)
  - logging [21-2](#)
- system processes

- displaying [26-2](#)
- status [26-5](#)
- system statistics
  - CPU and memory [26-5](#)
- system statistics reset feature [26-8](#)
- system switchover
  - configuring [4-4](#)
  - guidelines [4-4](#)
  - mechanisms [4-3](#)
- system switchovers
  - performing [5-20](#)
- SYSTEM variable [5-18](#)
  - clearing [5-18](#)

---

## T

- target disks [22-3](#)
- TCP connections
  - FCIP profiles [17-19](#)
- TCP statistics
  - displaying [17-10](#)
- Telnet
  - default service [14-18](#)
  - session [5-12, 5-20](#)
- Telnet session
  - message logging [21-5](#)
- temporary storage [2-12](#)
- TE port
  - trunking [1-4](#)
- TE ports
  - classes of service [9-4](#)
  - fctrace [24-3](#)
  - FSPF topology [15-2](#)
  - interface modes [9-2](#)
  - interoperability [24-20](#)
  - recovering from isolation [12-10](#)
  - SPAN [23-3](#)
  - trunking restrictions [10-1](#)
- TFTP

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- boot [5-28](#)
- copying images [5-12](#)
- server [5-28](#)
- TFTP server [26-6](#)
- threshold
  - major and minor [7-9](#)
- time interval
  - configuring [15-5](#)
- time out value
  - See TOV
- Timers
  - range [2-20](#)
- TL Ports
  - logging facility [21-3](#)
- TL ports
  - classes of service [9-3](#)
  - configuring [9-9](#)
  - displaying [9-22](#)
  - FCS [25-2, 25-3](#)
  - interface modes [9-2](#)
  - SPAN [23-3](#)
- TOV
  - interoperability [24-20](#)
  - ranges [24-2](#)
- troubleshooting
  - error messages [21-2](#)
- trunk-allowed list
  - configuring [10-4](#)
- Trunking
  - PortChannels comparison [11-3](#)
- trunking
  - configuration guidelines [10-6](#)
  - functionality [1-4](#)
  - interoperability [24-20](#)
  - link state [10-3](#)
  - restrictions [10-1](#)
- trunking ports [8-6](#)
- trunking protocol [10-2, 10-6](#)
  - default [10-2](#)
  - default settings [10-8](#)
- trunk mode
  - administrative default [9-13](#)
  - configuring [10-3](#)
  - default settings [10-8](#)
  - status [10-3](#)
- trunk ports
  - displaying information [10-7](#)

---

## U

- upgrade-reset feature [26-9](#)
- upgrades
  - See disruptive upgrades
  - See nondisruptive upgrades
- upgrading
  - software [5-12 to 5-24](#)
- upgrading BIOS
  - See BIOS upgrades
- user ID
  - authentication [14-2](#)
  - authorization process [14-4](#)
- user IDs
  - security management [1-8](#)
- user roles [1-8](#)
- users
  - creating [14-23](#)

---

## V

- VE ports [17-18](#)
- version compatibility
  - switch images [4-3](#)
- virtual devices [9-23](#)
- virtual E ports
  - See VE ports [17-18](#)
- virtual Fibre Channel hosts
  - mapping [17-46](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

virtual ISL [17-18](#)

Virtual LANs

See VLANs

virtual N port

dynamic mapping [17-46](#)

Virtual Router Redundancy Protocol

See VRRP

virtual SANs

See VSANs

VLANs

configuring [17-5](#)

VLAN tags [17-6](#)

VR IDs

configuring [16-14](#)

mapping [16-12](#)

VRRP

characteristics [16-12](#)

clearing statistics [16-18](#)

configuring [17-14](#)

configuring Gigabit Ethernet [17-13](#)

group members [17-13](#)

logging facility [21-3](#)

master and backup [16-13](#)

primary IP [16-14](#)

priority tracking [16-17](#)

security authentication [16-16](#)

setting priority [16-15](#)

tracking priority [16-16](#)

VSA

communicating attributes [14-16](#)

protocol options [14-17](#)

VSAN

address format [2-20](#)

configuring [1-4](#)

domain IDs [19-6](#)

functionality [1-4](#)

gateway switch [16-3](#)

overlaid routes [16-4](#)

reason codes [9-6](#)

redundancy [1-4](#)

scalability [1-4](#)

traffic isolation [1-4](#)

VSAN IDs

allowed list [10-8](#)

attributes [8-7](#)

FCS registration [2-8](#)

membership [8-4](#)

multiplexing traffic [9-4](#)

name [8-11](#)

range [8-6](#)

trunking [11-3](#)

VSANs

allowed-active [10-1, 10-4](#)

allowed list [23-3](#)

allowed-list [10-8](#)

attributes [8-7, 8-9](#)

availability [8-1](#)

broadcast address [15-9](#)

cache contents [19-14](#)

configuring [8-1, 8-7](#)

configuring domains [19-1](#)

configuring FSPF [15-3](#)

configuring overlay [16-10](#)

database submode [2-8](#)

default setting [8-11](#)

default VSAN [8-6](#)

deleting [8-9](#)

FCC protocol [20-2](#)

FCIDs [8-2](#)

FCS [25-2](#)

features [8-2](#)

flow statistics [15-12](#)

FSPF connectivity [15-2](#)

functionality [1-4](#)

interface [9-13, 9-15](#)

interop mode [24-20](#)

IP addresses [16-5](#)

IPFC interface [24-3](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- isolated VSAN [8-6](#)
- logical interface [3-9](#)
- loop devices [9-22](#)
- management interfaces [16-2](#)
- managing traffic [20-1](#)
- membership [8-6, 8-10](#)
- merging traffic [10-6](#)
- mismatch [9-7, 10-2](#)
- multiple zones [8-4, 12-6](#)
- name [8-7](#)
- name server [13-3](#)
- overlaid routes [16-8](#)
- port granularity [8-3](#)
- port isolation [10-6](#)
- Rules and features [14-7](#)
- sate [8-7](#)
- scalability [8-1](#)
- SPAN source [23-2, 23-3](#)
- static routing [16-6](#)
- TOVs [24-2](#)
- traffic isolation [8-1, 8-3](#)
- traffic routing [16-1](#)
- trunk allowed [9-13](#)
- trunk-allowed [10-1, 10-2](#)
- trunk-allowed list [10-4](#)
- trunking port [9-4](#)
- trunking ports [8-6](#)
- usage [8-10](#)
- VRRP [16-12](#)
- VRRP submode [2-8](#)

VSAN trunking

See trunking

watchdog checks [26-8](#)

world wide names

See WWNs

WWN

address pool [17-46](#)

WWNs

configuring [24-16](#)

displaying configurations [24-17](#)

suspended connection [9-7](#)

See also nWWNs

See also pWWNs

---

## Z

zone database [12-10](#)

zones

access control [12-6](#)

accesses between devices [1-4](#)

configuring [12-4](#)

configuring guidelines [12-6](#)

default policy [12-2, 12-9](#)

default settings [12-15](#)

enforcing [12-5](#)

examples [12-3](#)

functionality [1-4](#)

logging facility [21-3](#)

See also default zones

See also hard zoning

See also soft zoning

---

## W

warm switchovers

defining [4-4](#)

description [4-3](#)

guidelines [4-4](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***