



IP Services Configuration Guide, Cisco DCNM for SAN

Cisco DCNM for SAN, Release 6.x

October 2014

Cisco Systems, Inc.

www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Text Part Number: OL-27520-01

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)

IP Services Configuration Guide, Cisco DCNM for SAN
© 2012—2014 Cisco Systems, Inc. All rights reserved.



Preface	xv
Audience	xv
Organization	xv
Document Conventions	xvi
Related Documentation	xvi
Release Notes	xvii
Regulatory Compliance and Safety Information	xvii
Compatibility Information	xvii
Hardware Installation	xvii
Software Installation and Upgrade	xvii
Cisco NX-OS	xvii
Cisco DCNM-SAN	xviii
Command-Line Interface	xviii
Intelligent Storage Networking Services Configuration Guides	xviii
Troubleshooting and Reference	xviii
Obtaining Documentation and Submitting a Service Request	xix
New and Changed Information	1
1	
IP Services Overview	1
FCIP	1
SAN Extension Tuner	2
iSCSI	2
IP Services	2
IP Storage	2
IPv4 and IPv6	2
1	
Configuring FCIP	1
Information About FCIP	1
FCIP Concepts	2
FCIP and VE Ports	2
FCIP Links	3
FCIP Profiles	4
FCIP Interfaces	5
FCIP High-Availability Solutions	5
Fibre Channel PortChannels	5
FSPF	6
VRRP	7
Ethernet PortChannels	7
Ethernet PortChannels and Fibre Channel PortChannels	8
FCIP Profile Configuration	8
Text Part Number:	

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Peers	9
FCIP B Port Interoperability Mode	9
Quality of Service	11
E Ports	12
FCIP Write Acceleration	12
FCIP Tape Acceleration	14
FCIP Compression	18
Default Settings	19
Configuring FCIP	20
Enabling FCIP	21
Modifying an FCIP Link	23
Creating FCIP Profiles	23
Checking Trunk Status	24
Launching Cisco Transport Controller	24
Creating FCIP Links	25
Configuring TCP Listener Ports	25
Configuring TCP Parameters	26
Assigning a Peer IP Address	31
Configuring Active Connections	32
Enabling Time Stamp Control	33
Configuring B Ports	34
Setting QoS Values	35
Configuring FCIP Write Acceleration	35
Configuring FCIP Tape Acceleration	36
Configuring FCIP Compression	37
Verifying FCIP Configuration	37
Displaying FCIP Profile Information	38
Displaying FCIP Profile Configuration Information	39
Displaying FCIP Profile Configuration Information	39
Displaying FCIP Interface Information	39
Displaying Write Acceleration Activity Information	41
Displaying Tape Acceleration Activity Information	42
Displaying FCIP Compression Information	44
Field Descriptions for FCIP	45
FCIP Monitor	45
FCIP Interfaces	45
FCIP Interfaces Trunk Failures	46
FCIP FICON Configuration	46
FCIP Profiles	46
FCIP Tunnels	47
FCIP Tunnels (Advanced)	48
FCIP Tunnels (FICON TA)	49
FCIP Tunnels Statistics	49
FCIP XRC Statistics	49
Additional References	49

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Related Document	50
Standards	50
RFCs	50
MIBs	50
Feature History for FCIP	50
1	
Configuring the SAN Extension Tuner	1
Information About the SAN Extension Tuner	1
SAN Extension Tuner Setup	3
Data Pattern	4
Licensing Requirements for SAN Extension Tuner	4
Default Settings	4
Configuring the SAN Extension Tuner	5
Tuning the FCIP Link	5
Using the SAN Extension Tuner Wizard	5
Enabling the Tuner	7
Configuring nWWN	7
Configuring the Virtual N Port	7
Assigning the SCSI Read/Write	8
Assigning SCSI Tape Read/Write	10
Configuring a Data Pattern	11
Verifying SAN Extension Tuner Configuration	12
Verifying the SAN Extension Tuner Configuration	12
Additional References	13
Related Document	14
Standards	14
RFCs	14
MIBs	14
1	
Configuring iSCSI	1
Information About iSCSI	2
About iSCSI Configuration Limits	4
Presenting Fibre Channel Targets as iSCSI Targets	4
Dynamic Mapping	5
Static Mapping	6
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts	6
Initiator Identification	6
Initiator Presentation Modes	7
Transparent Initiator Mode	7
WWN Assignment for iSCSI Initiators	8
Static Mapping	9
Proxy Initiator Mode	9
VSAN Membership for iSCSI	11
Advanced VSAN Membership for iSCSI Hosts	11
iSCSI Access Control	11

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Fibre Channel Zoning-Based Access Control	11
iSCSI-Based Access Control	12
Enforcing Access Control	13
iSCSI Session Authentication	13
iSCSI Immediate Data and Unsolicited Data Features	14
iSCSI Listener Port	14
TCP Tuning Parameters	14
iSCSI Routing Modes	15
About iSLB	17
About iSLB Initiators	17
Assigning WWNs to iSLB Initiators	18
iSLB Initiator Targets	18
iSLB Session Authentication	19
About Load Balancing Using VRRP	19
Changing iSCSI Interface Parameters and the Impact on Load Balancing	21
VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces	21
About iSLB Configuration Distribution Using CFS	21
Locking the Fabric	22
CFS Merge Process	22
iSLB CFS Merge Status Conflicts	23
iSCSI High Availability	23
Transparent Target Failover	23
iSCSI High Availability with Host Running Multi-Path Software	23
iSCSI HA with Host Not Having Any Multi-Path Software	24
LUN Trespass for Storage Port Failover	25
Multiple IPS Ports Connected to the Same IP Network	26
VRRP-Based High Availability	27
Ethernet PortChannel-Based High Availability	28
iSNS	29
About iSNS Client Functionality	30
About iSNS Server Functionality	30
iSNS Client Registration and Deregistration	31
Target Discovery	31
About Cloud Discovery	31
Licensing Requirements for iSCSI	32
Guidelines and Limitations	32
Default Settings	33
Configuring iSCSI	34
Enabling iSCSI	35
Creating iSCSI Interfaces	36
Using the iSCSI Wizard	37
Enabling Dynamic Mapping	37
Creating Static Mapping	38
Advertising Static iSCSI Targets	39
Specifying the Initiator Identification	39

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Configuring the iSCSI Initiator Idle Timeout	40
Configuring Dynamic Mapping	41
Configuring Static Mapping	41
Making the Dynamic iSCSI Initiator WWN Mapping Static	43
Checking for WWN Conflicts	43
Configuring the Proxy Initiator	44
Configuring VSAN Membership for iSCSI Hosts	45
Configuring Default Port VSAN for iSCSI Interfaces	46
Adding iSCSI Initiator to the Zone Database	47
Configuring Access Control in iSCSI	48
Configuring AAA Authentication for an iSCSI User	50
Configuring Authentication Mechanism	50
Configuring Local Authentication	52
Restricting iSCSI Initiator Authentication	52
Configuring Mutual CHAP Authentication	53
Configuring an iSCSI RADIUS Server	54
Setting QoS Values	55
Changing the iSCSI Routing Mode	55
Configuring iSLB	56
Configuring iSLB Using Device Manager	56
Configuring iSLB Initiator Names or IP Addresses	57
Making the Dynamic iSLB Initiator WWN Mapping Static	59
Assigning VSAN Membership for iSLB Initiators	59
Configuring Metric for Load Balancing	60
Configuring iSLB Initiator Targets	61
Configuring and Activating Zones for iSLB Initiators and Initiator Targets	63
Restricting iSLB Initiator Authentication	64
Mutual CHAP Authentication	65
Configuring Load Balancing Using VRRP	66
Enabling VRRP for Load Balancing	66
Distributing the iSLB Configuration Using CFS	67
Enabling iSLB Configuration Distribution	67
Committing Changes to the Fabric	68
Discarding Pending Changes	68
Clearing a Fabric Lock	69
Creating a Static iSCSI Virtual Target	69
Enabling the Trespass Feature for a Static iSCSI	71
Configuring iSCSI Authentication	71
Configuring No Authentication	72
Configuring CHAP with Local Password Database	72
Configuring CHAP with External RADIUS Server	73
Creating an iSNS Client Profile	75
Configuring iSNS Servers	77
Enabling the iSNS Server	78
iSNS Configuration Distribution	78

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Configuring the ESI Retry Count	79
Configuring the Registration Period	80
Configuring iSNS Cloud Discovery	80
Enabling iSNS Cloud Discovery	81
Initiating On-Demand iSNS Cloud Discovery	81
Configuring Automatic iSNS Cloud Discovery	82
Configuring iSNS Cloud Discovery Distribution	82
Configuring iSNS Cloud Discovery Message Types	83
Verifying iSCSI Configuration	83
Displaying iSCSI Interfaces	86
Displaying iSCSI Statistics	87
Displaying Proxy Initiator Information	88
Displaying Global iSCSI Information	90
Displaying iSCSI Sessions	90
Displaying iSCSI Initiators	91
Displaying iSCSI Virtual Targets	95
Displaying iSCSI User Information	95
Displaying iSLB VRRP Information	95
Displaying Pending iSLB Configuration Changes	96
Displaying iSLB CFS Status	96
Displaying iSLB CFS Distribution Session Status	96
Displaying iSLB CFS Merge Status	96
Verifying iSNS Client Configuration	97
Verifying the iSNS Server Configuration	98
Verifying Automatic iSNS Cloud Discovery Configuration	104
Verifying Cloud Discovery Status	105
Verifying Cloud Discovery Membership	105
Displaying Cloud Discovery Statistics	105
Configuration Examples for iSCSI	105
Example of VSAN Membership for iSCSI Devices	107
Example of an iSNS Server	110
iSCSI Transparent Mode Initiator Example	111
Target Storage Device Requiring LUN Mapping Example	119
Field Descriptions for iSCSI	128
Ethernet Interfaces iSCSI	129
Ethernet Interfaces iSCSI TCP	129
Ethernet Interface Monitor iSCSI Connections	130
iSCSI Connection	130
iSCSI Initiators	131
iSCSI Targets	132
iSCSI Session Initiators	132
iSCSI Global	133
iSCSI Session Statistics	133
iSCSI iSLB VRRP	133
iSCSI Initiator Access	134

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

iSCSI Initiator PWWN	134
iSCSI Sessions	134
iSCSI Sessions Detail	134
iSNS Details iSCSI Nodes	135
iSCSI User	135
Edit iSCSI Advertised Interfaces	135
Additional References	136
Related Document	136
Standards	136
RFCs	136
MIBs	136
1	
Configuring IP Services	1
Information About IP Services	1
Traffic Management Services	2
Management Interface Configuration	3
About the Default Gateway	3
IPv4 Default Network Configuration	4
IPFC	5
About IPv4 Static Routes	5
About Overlay VSANs	5
About VRRP	5
DNS Server Configuration	7
Guidelines and Limitations	7
Default Settings	8
Configuring IP Services	8
Configuring Management Interface	9
Configuring the Default Gateway	10
Configuring Default Networks using IPV4	11
Configuring an IPv4 Address in a VSAN	11
Enabling IPv4 Routing	11
Configuring IPv4 Static Routes	12
Configuring Overlay VSANs	12
Configuring Multiple VSANs	14
Configuring VRRP	16
Adding and Deleting a Virtual Router	17
Virtual Router Initiation	17
Adding Virtual Router IP Addresses	18
Setting the Priority for the Virtual Router	19
Setting the Time Interval for Advertisement Packets	20
Configuring or Enabling Priority Preemption	21
Setting Virtual Router Authentication	22
Tracking the Interface Priority	23
Configuring DNS Server	24
Verifying IP Services Configuration	25

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- Verifying the Default Gateway Configuration 25
- Verifying the VSAN Interface Configuration 26
- Verifying the IPv4 Routing Configuration 26
- Verifying IPv4 Static Route Information 26
- Displaying and Clearing ARPs 27
- Displaying IPv4 VRRP Information 27
- Displaying IPv6 VRRP Information 28
- Displaying VRRP Statistics 29
- Clearing VRRP Statistics 29
- Displaying DNS Host Information 30
- Configuration Examples for IP Services 30
- Field Descriptions for IP Services 33
- IP Routes 33
- IP Statistics ICMP 33
- IP Statistics IP 34
- IP Statistics SNMP 35
- IP Statistics UDP 37
- mgmt0 Statistics 37
- TCP UDP TCP 37
- TCP UDP UDP 37
- VRRP General 38
- VRRP IP Addresses 38
- VRRP Statistics 39
- CDP General 39
- CDP Neighbors 40
- iSNS Profiles 40
- iSNS Servers 40
- iSNS Entities 41
- iSNS Cloud Discovery 41
- iSNS Clouds 41
- iSNS Cloud Interfaces 42
- Monitor Dialog Controls 42
- iSNS Details iSCSI Nodes 43
- iSNS Details Portals 43
- Additional References 43
- Related Document 44
- Standards 44
- RFCs 44
- MIBs 44
- 1
- Configuring IP Storage 1
- Information About IP Storage 1
- IPS Module Upgrade 3
- MPS-14/2 Module Upgrade 4
- Supported Hardware 4

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Gigabit Ethernet Interfaces for IPv4 Configuration	4
Basic Gigabit Ethernet Configuration	5
IPS Module Core Dumps	5
Interface Descriptions Configuration	6
Beacon Mode Configuration	6
Autonegotiation Configuration	6
MTU Frame Size Configuration	6
Promiscuous Mode Configuration	6
About VLANs for Gigabit Ethernet	6
Interface Subnet Requirements	7
Verifying Gigabit Ethernet Connectivity	7
Gigabit Ethernet High Availability	8
VRRP for iSCSI and FCIP Services	8
About Ethernet PortChannel Aggregation	9
CDP	10
Licensing Requirements for IP Storage	10
Guidelines and Limitations	10
Default Settings	11
Configuring IP Storage	11
Configuring IPS Core Dumps	11
Configuring VRRP for Gigabit Ethernet Interfaces	12
Configuring Ethernet PortChannels	14
Verifying IP Storage Configuration	15
Verifying Module Status	15
Displaying Gigabit Ethernet Interface Statistics	16
Displaying Ethernet MAC Statistics	17
Displaying DMA-Bridge Statistics	17
Displaying TCP Statistics	18
Field Descriptions for IP Storage	20
FCIP Profiles	20
FCIP Tunnels	21
FCIP Tunnels (Advanced)	21
FCIP Tunnels (FICON TA)	22
FCIP Tunnels Statistics	22
FCIP XRC Statistics	22
iSCSI Connection	23
iSCSI Initiators	23
iSCSI Targets	24
iSCSI Session Initiators	25
Module Control	25
iSCSI Global	25
iSCSI Session Statistics	26
iSCSI iSLB VRRP	26
iSCSI Initiator Access	26
Initiator Specific Target	27


[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- iSCSI Initiator PWWN 27
- iSCSI Sessions 28
- iSCSI Sessions Detail 28
- Additional References 28
- Related Document 29
- Standards 29
- RFCs 29
- MIBs 29
- 1
- Configuring IPv4 for Gigabit Ethernet Interfaces 1
- Information About IPv4 1
- Interface Descriptions 2
- Beacon Mode 2
- About VLANs for Gigabit Ethernet 2
- Interface Subnet Requirements 3
- Licensing Requirements for IPv4 for Gigabit Ethernet Interfaces 3
- Guidelines and Limitations 4
- Default Settings 4
- Configuring IPv4 4
- Configuring Gigabit Ethernet Interface 5
- Configuring Autonegotiation 5
- Configuring the MTU Frame Size 6
- Configuring Promiscuous Mode 7
- Configuring the VLAN Subinterface 7
- Configuring Static IPv4 Routing 8
- Applying IPv4-ACLs on Gigabit Ethernet Interfaces 8
- Clearing ARP Cache 9
- Verifying IPV4 Configuration 9
- Verifying Gigabit Ethernet Connectivity 10
- Displaying the IPv4 Route Table 10
- Displaying ARP Cache 11
- Displaying IPv4 Statistics 11
- Configuration Examples for IPV4 11
- Additional References 12
- Related Document 13
- Standards 13
- RFCs 13
- MIBs 13
- 1
- Configuring IPv6 for Gigabit Ethernet Interfaces 1
- Information About IPV6 1
- Extended IPv6 Address Space for Unique Addresses 2
- IPv6 Address Formats 2
- IPv6 Address Prefix Format 3
- IPv6 Address Type: Unicast 3

Send documentation comments to dcnm-san-docfeedback@cisco.com

Global Addresses	3
Link-Local Address	4
IPv6 Address Type: Multicast	5
ICMP for IPv6	6
Path MTU Discovery for IPv6	7
IPv6 Neighbor Discovery	7
IPv6 Neighbor Solicitation and Advertisement Messages	7
Router Discovery	9
IPv6 Stateless Autoconfiguration	9
Dual IPv4 and IPv6 Protocol Stacks	10
IPv6 Addressing and Enabling IPv6 Routing	11
Transitioning from IPv4 to IPv6	12
Guidelines and Limitations	12
Default Settings	13
Configuring Basic Connectivity for IPv6	13
Configuring IPv6 Addressing and Enabling IPv6 Routing	14
Configuring IPv4 and IPv6 Protocol Addresses	15
Clearing IPv6 Neighbor Discovery Cache	16
Configuring Neighbor Discovery Parameters	16
Duplicate Address Detection Attempts	16
Reachability Time	17
Retransmission Time	17
Configuring a IPv6 Static Route	17
Verifying IPV6 Configuration	18
Verifying Neighbor Discovery Parameter Configuration	19
Verifying IPv6 Static Route Configuration and Operation	19
Displaying IPv6	20
Configuration Examples for IPV6	21
Example Output for the show ipv6 interface Command	21
Example Output for the show ipv6 neighbours Command	22
Example Output for the show ipv6 traffic Command	22
Additional References	23
Related Document	23
Standards	23
RFCs	23
MIBs	23
i	

Send documentation comments to dcnm-san-docfeedback@cisco.com



Preface

This preface describes the audience, organization, and conventions of the *IP Services Configuration Guide, Cisco DCNM for SAN*. 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	IP Services Overview	Provides an overview of the Intelligent Storage Services supported by the Cisco MDS 9000 NX-OS software.
Chapter 2	Configuring FCIP	Describes how the switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.
Chapter 3	Configuring the SAN Extension Tuner	Explains the SAN extension tuner (SET) feature that optimizes FCIP performance.
Chapter 4	Configuring iSCSI	Describes the iSCSI feature that is specific to the IPS module and is available in the Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.
Chapter 5	Configuring IP Services	Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.



Chapter	Title	Description
Chapter 6	Configuring IP Storage	Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together through IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol.
Chapter 7	Configuring IPv4 for Gigabit Ethernet Interfaces	Describes the IPv4 protocol support provided by Cisco MDS 9000 Family switches.
Chapter 8	Configuring IPv6 for Gigabit Ethernet Interfaces	Describes the IPv6 protocol support provided by Cisco MDS 9000 Family switches.

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.

This document uses the following conventions:

- Note** Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.
- Caution** Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocator.htm

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS NX-OS Releases*
- *Cisco MDS 9000 Family Release Notes for MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- Cisco DCNM Release Notes

Regulatory Compliance and Safety Information

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

Compatibility Information

- Cisco Data Center Interoperability Support Matrix
- Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

Hardware Installation

- Cisco MDS 9500 Series Hardware Installation Guide
- Cisco MDS 9200 Series Hardware Installation Guide
- *Cisco MDS 9100 Series Hardware Installation Guide*
- Cisco MDS 9124 and Cisco MDS 9134 Multilayer Fabric Switch Quick Start Guide

Software Installation and Upgrade

- Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide

Cisco NX-OS

- Cisco MDS 9000 Family NX-OS Licensing Guide
- Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide
- Cisco MDS 9000 Family NX-OS System Management Configuration Guide
- Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide
- Cisco MDS 9000 Family NX-OS Fabric Configuration Guide
- Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide
- Cisco MDS 9000 Family NX-OS Security Configuration Guide
- Cisco MDS 9000 Family NX-OS IP Services Configuration Guide
- *Cisco MDS 9000 Family NX-OS Intelligent Storage Services Configuration Guide*

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide](#)
- [Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide](#)
- [Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS](#)

Cisco DCNM-SAN

- [Cisco DCNM Fundamentals Guide, Release 6.x](#)
- [System Management Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Interfaces Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Fabric Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Quality of Service Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Security Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [IP Services Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Intelligent Storage Services Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [Inter-VSAN Routing Configuration Guide, Cisco DCNM for SAN, Release 6.x](#)
- [SMI-S and Web Services Programming Guide, Cisco DCNM for SAN, Release 6.x](#)

Command-Line Interface

- [Cisco MDS 9000 Family Command Reference](#)

Intelligent Storage Networking Services Configuration Guides

- [Cisco MDS 9000 Family I/O Acceleration Configuration Guide](#)
- [Cisco MDS 9000 Family SANTap Deployment Guide](#)
- [Cisco MDS 9000 Family Data Mobility Manager Configuration Guide](#)
- [Cisco MDS 9000 Family Storage Media Encryption Configuration Guide](#)

Troubleshooting and Reference

- [Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference](#)
- [Cisco MDS 9000 Family SAN-OS Troubleshooting Guide](#)
- [Cisco MDS 9000 Family NX-OS MIB Quick Reference](#)
- [Cisco DCNM for SAN Database Schema Reference](#)

Obtaining Documentation and Submitting a Service Request

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

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

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

Send documentation comments to dcnm-san-docfeedback@cisco.com



New and Changed Information

As of Cisco DCNM Release 6.x, Cisco Fabric Manager and Cisco Data Center Network Manager for LAN are merged into one unified product called Cisco Data Center Network Manager (DCNM) that can manage both LAN and SAN environments. As a part of this product merger, the name Cisco DCNM for SAN replaces the name Cisco Fabric Manager.

The following documentation changes support the merged Cisco DCNM product:

- Cisco DCNM product documentation for Cisco DCNM Release 6.x is retitled with the name Cisco DCNM for LAN.
- Cisco Fabric Manager product documentation for Cisco DCNM Release 6.x is retitled with the name Cisco DCNM for SAN.
- Cisco DCNM for SAN product documentation is now published to the Data Center Network Manager listing page on Cisco.com: http://www.cisco.com/en/US/products/ps9369/tsd_products_support_configure.html

This URL is also the listing page for Cisco DCNM for LAN product documentation.

- Cisco Fabric Manager documentation for software releases earlier than Cisco DCNM Release 6.x, retains the name Cisco Fabric Manager and remains available at its current Cisco.com listing page: http://www.cisco.com/en/US/products/ps10495/tsd_products_support_configure.html

You should continue to use the Cisco Fabric Manager documentation if you are using a release of Cisco Fabric Manager software that is earlier than Cisco DCNM Release 6.x.

- The name DCNM-SAN is used in place of Cisco DCNM for SAN in the user interface of Cisco Data Center Network Manager; likewise, the name DCNM-LAN is used in place of Cisco DCNM for LAN in the user interface. To match the user interface, the product documentation also uses the names DCNM-SAN and DCNM-LAN.
- The following new publications support both Cisco DCNM for LAN and DCNM for SAN, and address the new licensing model, the new installation process, and the new features of Cisco DCNM:
 - *Cisco DCNM Installation and Licensing Guide*
 - Cisco DCNM Release Notes
- For a complete list of Cisco DCNM documentation, see the “Related Documentation” section in the Preface.

About This Guide

The information in the new *Cisco Fabric Manager IP Services Configuration Guide* previously existed in Part 6: IP Services of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

The following table lists the New and Changed features for this guide, starting with Release 4.2(1):



Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1 ***New and Changed Features***

Feature	New or Changed Topics	Changed in Release	Where Documented
All occurrences for Fabric Manager changed to DCNM for SAN	DCNM for SAN	5.2(1)	New and Changed Information
Configuring FCIP	FCIP Compression	5.0(1a)	Chapter 2, “Configuring FCIP”



IP Services Overview

The Cisco MDS 9000 NX-OS software provides features such as FCIP, SAN Extension Tuner, iSCSI, IP storage, IPv4, and IPv6 in a single platform. These IP services simplify SAN provisioning by automatically distributing configuration information to all the switches in a storage network. The Virtual Routing Redundancy Protocol (VRRP) increases the IP network availability for iSCSI and FCIP connections by allowing failover of connections from one port to another. The increased IP network availability facilitates the failover of an iSCSI volume from one IP services port to any other IP services port, either locally or on another Cisco MDS 9000 switch.

This chapter includes the following sections:

- [FCIP, page 1-1](#)
- [SAN Extension Tuner, page 1-2](#)
- [iSCSI, page 1-2](#)
- [IP Services, page 1-2](#)
- [IP Storage, page 1-2](#)
- [IPv4 and IPv6, page 1-2](#)

FCIP

FCIP (Fibre Channel over IP Protocol) transparently connects a remote Fibre Channel storage area network (SAN island) by transporting Fibre Channel data from a local SAN to a remote SAN using IP networks. IP network availability for the FCIP connections can be increased by using features such as Virtual Routing Redundancy Protocol (VRRP) and quality of service (QoS). FCIP can be optimized for wire performance through enhancements that address out-of-order delivery issues, support jumbo frames, provide traffic shaping, and perform TCP optimization.

For more information on configuring FCIP, see [Chapter 2, “Configuring FCIP.”](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

SAN Extension Tuner

The SAN Extension Tuner (SET) feature helps you optimize FCIP performance by generating Small Computer System Interface (SCSI) I/O commands and directing the traffic to a specific virtual target. SET reports the I/Os per second and I/O latency results, which helps you to determine the number of concurrent I/Os needed to maximize the FCIP throughput.

For information on configuring the SAN Extension Tuner, see [Chapter 3, “Configuring the SAN Extension Tuner.”](#)

iSCSI

The iSCSI feature allows an IP host to access Fibre Channel storage. This feature enables routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN. The Fibre Channel Storage devices are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch.

For information on configuring iSCSI, see [Chapter 4, “Configuring iSCSI.”](#)

IP Services

The IP Services Modules allow you to extend storage networks using the Ethernet infrastructure. The Cisco MDS 9000 Family switches route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route the traffic between VSANs. This chapter also describes the procedure to configure IP Route using DCNM for SAN And Device Manager. From NX-OS release 4.2(1) and later, CPP interfaces are also available for selection while creating a new IP route.

For information on configuring IP services, see [Chapter 5, “Configuring IP Services.”](#)

IP Storage

The IP Storage (IPS) Service module allows you to use the open-standard FCIP protocol to enable interconnection of SAN islands over extended distances. The IPS module and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features that are available on other switching modules, including VSANs, security, and traffic management.

For information on configuring IP Storage, see [Chapter 6, “Configuring IP Storage.”](#)

IPv4 and IPv6

The Cisco MDS 9000 NX-OS software supports the IP version 4 (IPv4) and version 6 (IPv6) protocols on Gigabit Ethernet interfaces. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6, while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses. The dual stack approach for IPv4 and IPv6 allows Cisco MDS 9000 Family switches to connect to older IP networks, transitional networks of both versions, and IPv6 data networks.

Send documentation comments to dcnm-san-docfeedback@cisco.com

For more information on configuring IPv4 for Gigabit Ethernet interfaces, see [Chapter 7, “Configuring IPv4 for Gigabit Ethernet Interfaces.”](#)

For more information on configuring IPv6 for Gigabit Ethernet interfaces, see [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)



Configuring FCIP

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch can connect separated SAN islands using Fibre Channel over IP (FCIP).



Note

FCIP is supported on the MDS 9222i switch, MSM-18/4 module, MDS 9216i switch, MPS-14/2 module, 16-Port Storage Services Node (SSN-16), and IPS modules on MDS 9200 Series directors.


This chapter includes the following topics:

- [Information About FCIP, page 2-1](#)
- [Default Settings, page 2-19](#)
- [Configuring FCIP, page 2-20](#)
- [Verifying FCIP Configuration, page 2-37](#)
- [Field Descriptions for FCIP, page 2-45](#)
- [Additional References, page 2-49](#)
- [Feature History for FCIP, page 2-50](#)

Information 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). The switch can connect separated SAN islands using Fibre Channel over IP (FCIP) (see [Figure 2-1](#)).

Figure 2-1 *Fibre Channel SANs Connected by FCIP*



FCIP uses TCP as a network layer transport. The DF bit is set in the TCP header.

Text Part Number:

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

For more information about FCIP protocols, refer to the IETF standards for IP storage at <http://www.ietf.org>. Also refer to Fibre Channel standards for switch backbone connection at <http://www.t11.org> (see FC-BB-2).

This section includes the following topics:

- [FCIP Concepts, page 2-2](#)
- [FCIP High-Availability Solutions, page 2-5](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 2-8](#)
- [FCIP Profile Configuration, page 2-8](#)
- [Peers, page 2-9](#)
- [Quality of Service, page 2-11](#)
- [E Ports, page 2-12](#)
- [FCIP Write Acceleration, page 2-12](#)
- [FCIP Tape Acceleration, page 2-14](#)
- [FCIP Compression, page 2-18](#)

FCIP Concepts

To configure IPS modules or MPS-14/2 modules for FCIP, you should have a basic understanding of the following concepts:

- [FCIP and VE Ports, page 2-2](#)
- [FCIP Links, page 2-3](#)
- [FCIP Profiles, page 2-46](#)
- [FCIP Interfaces, page 2-5](#)

FCIP and VE Ports


Figure 2-2 describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's extended ISLs (EISLs).

FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see Figure 2-2).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-2 FCIP Links and Virtual ISLs



91557

See the “Configuring B Ports” section on page 2-34 for more information.

FCIP Links

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link:

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch protocol frames (all Class F). This arrangement provides low latency for all control frames.

To enable FCIP on the IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured.


The FCIP link is established between two peers, the VE port initialization behavior is identical to a 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.

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules or MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (see Figure 2-3).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-3 Assigning Profiles to Each Gigabit Ethernet Interface




FCIP Profiles

The FCIP profile contains information about the 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 terminate (see Figure 2-4).


Figure 2-4 FCIP Profile and FCIP Links



You must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile. You can assign IPv4 or IPv6 addresses to the interfaces. Figure 2-5 shows an example configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-5 Assigning Profiles to Each Gigabit Ethernet Interface



91561

FCIP Interfaces

The FCIP interface is the local endpoint 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 initiates the 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.

FCIP High-Availability Solutions

The following high-availability solutions are available for FCIP configurations:


- [Fibre Channel PortChannels, page 2-5](#)
- [FSPF, page 2-6](#)
- [VRRP, page 2-7](#)
- [Ethernet PortChannels, page 2-7](#)

Fibre Channel PortChannels

Figure 2-6 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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-6 *PortChannel-Based Load Balancing*



90857


The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.

FSPF

Figure 2-7 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 2-7 *FSPF-Based Load Balancing*



90858

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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

VRRP

Figure 2-8 displays a Virtual Router Redundancy Protocol (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 2-8 VRRP-Based High Availability




The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Ethernet PortChannels

Figure 2-9 displays an Ethernet PortChannel-based high-availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 2-9 Ethernet PortChannel-Based High Availability



The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link-level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet link.

Send documentation comments to dcnm-san-docfeedback@cisco.com


- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer link redundancy between the Cisco MDS 9000 Family switch's Gigabit Ethernet ports and the connecting Ethernet switch. Fibre Channel PortChannels also offer (E)ISL link redundancy between Fibre Channel switches. FCIP is an (E)ISL link and is only applicable for a Fibre Channel PortChannel. Beneath the FCIP level, an FCIP link can run on top of an Ethernet PortChannel or on one Gigabit Ethernet port. This link is totally transparent to the Fibre Channel layer.

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet PortChannel (see [Chapter 6, “Configuring IP Storage”](#) for more information). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortChannel (to which FCIP link can be a part of) does not have a restriction on which (E)ISL links can be combined in a Fibre Channel PortChannel as long as it passes the compatibility check. The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see [Figure 2-10](#)).

Figure 2-10 PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* [Interfaces Configuration Guide, Cisco DCNM for SAN](#).

To configure Ethernet PortChannels, see the *High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN* [Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide](#).

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.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

FCIP configuration options can be accessed from the switch(Config-profile)# submode prompt.

Peers

All the FCIP and E port parameters are configured in context to the FCIP interface. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch. The basic FCIP configuration uses the peer's IP address to configure the peer information. You can establish an FCIP link with the peer using the Peer IP address option. This option configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.

To establish an FCIP link with the peer, you can use the peer IP address option. This option configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.

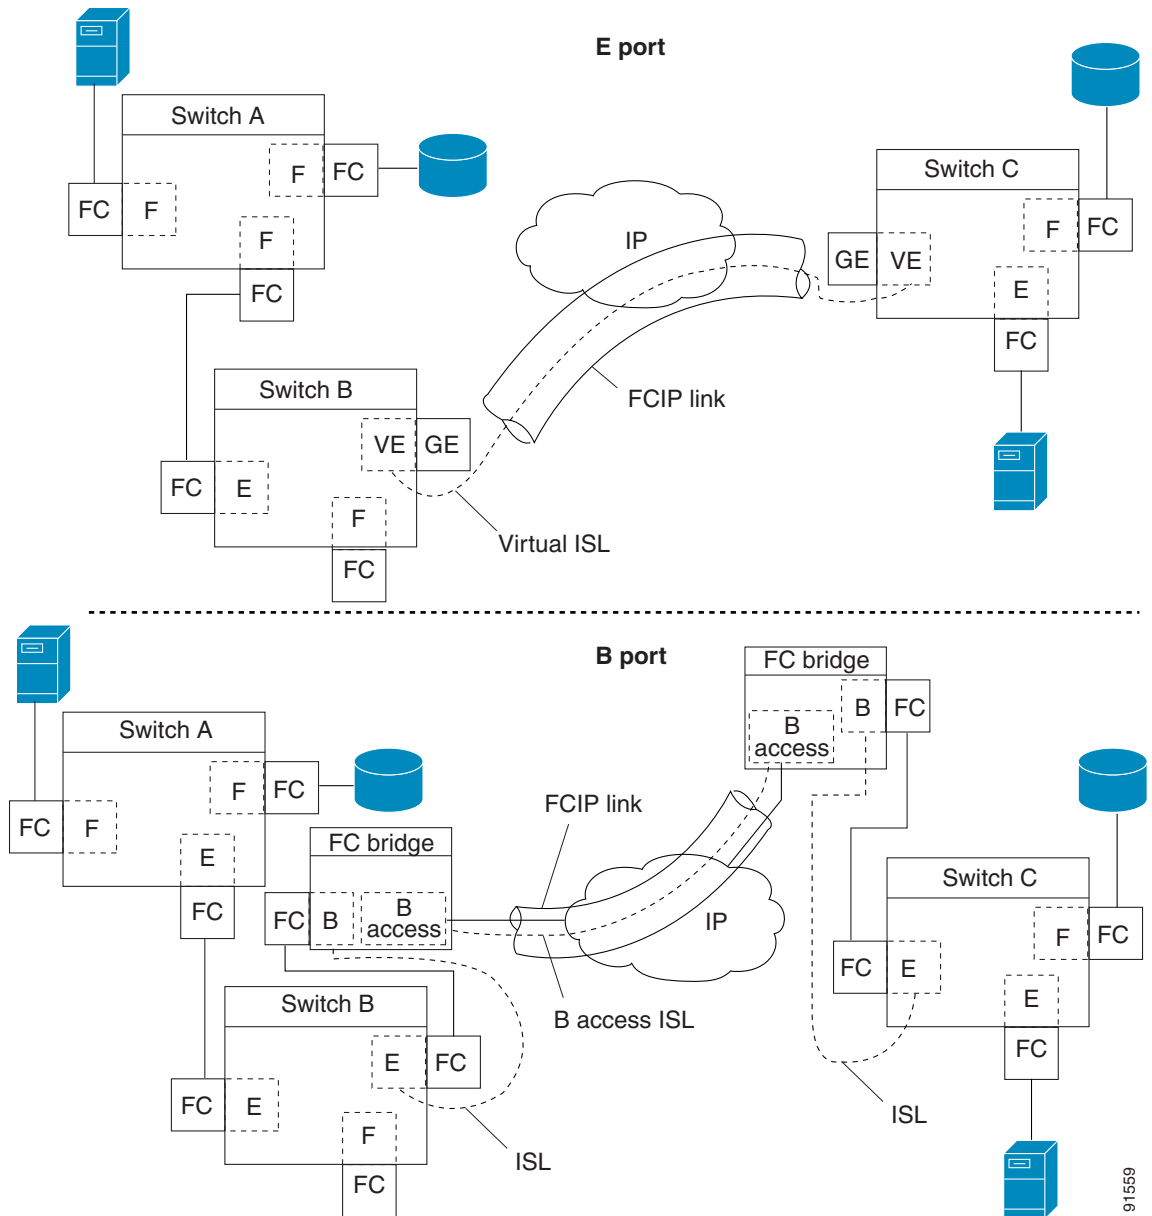
To establish a peer connection, you must first create the FCIP interface and enter the config-if submode.

FCIP 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 2-11](#) shows a typical SAN extension over an IP network.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-11 FCIP B Port and Fibre Channel E Port



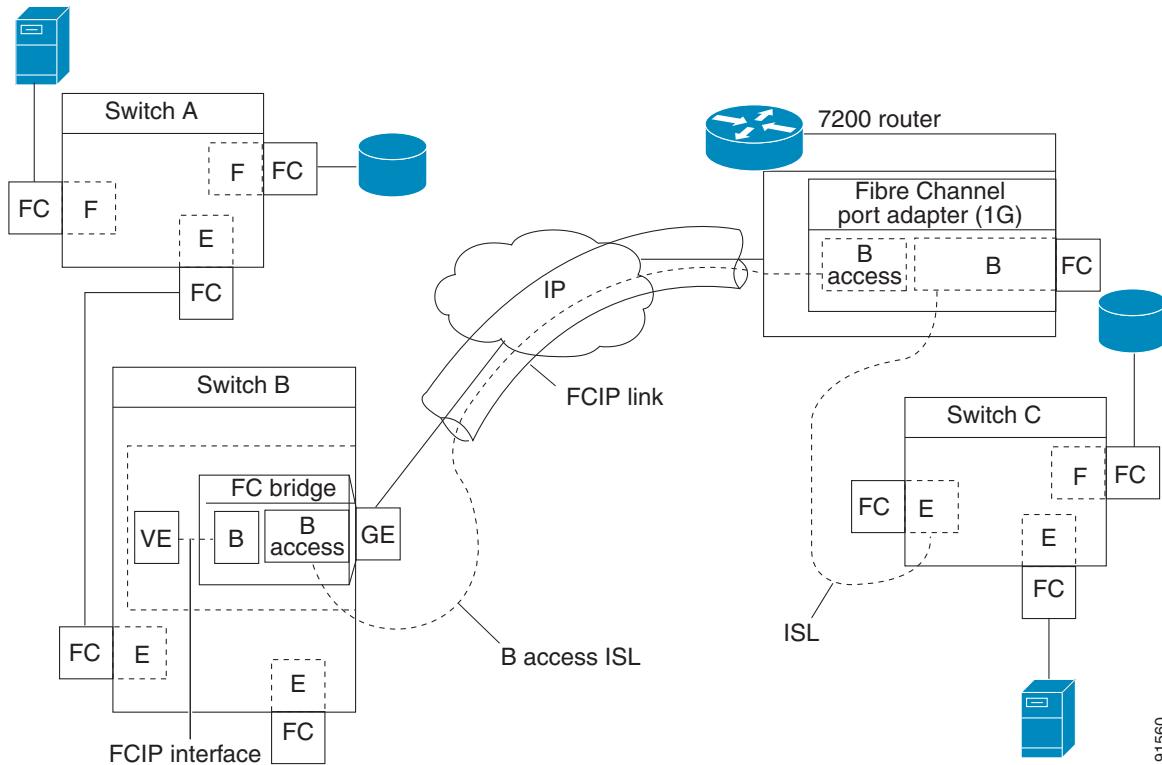
B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over an FCIP link, B ports use a B access ISL.*

Send documentation comments to dcnm-san-docfeedback@cisco.com

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see Figure 2-12).

Figure 2-12 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, eliminating the need for local bridge devices.

Quality of Service

The quality of service (QoS) parameter specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that connection.

Send documentation comments to dcnm-san-docfeedback@cisco.com

E Ports

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN
- Trunk mode and trunk allowed VSANs
- PortChannels
- FSPF
- Fibre Channel domains (fcdomains)
- Importing and exporting the zone database from the adjacent switch

You can configure E ports in the same way you configure FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN
See the *Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.
- Trunk mode and trunk allowed VSANs
See the *Interfaces Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*.
- PortChannels
 - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
 - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
 See the *Security Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- FSPF
See the *Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.
- Fibre Channel domains (fcdomains)
See the *System Management Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Importing and exporting the zone database from the adjacent switch
See the *System Management Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

FCIP Write Acceleration

The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Note

The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel the write acceleration feature will be turned operationally off.



Note

IBM Peer to Peer Remote Copy (PPRC) is not supported with FCIP write acceleration.

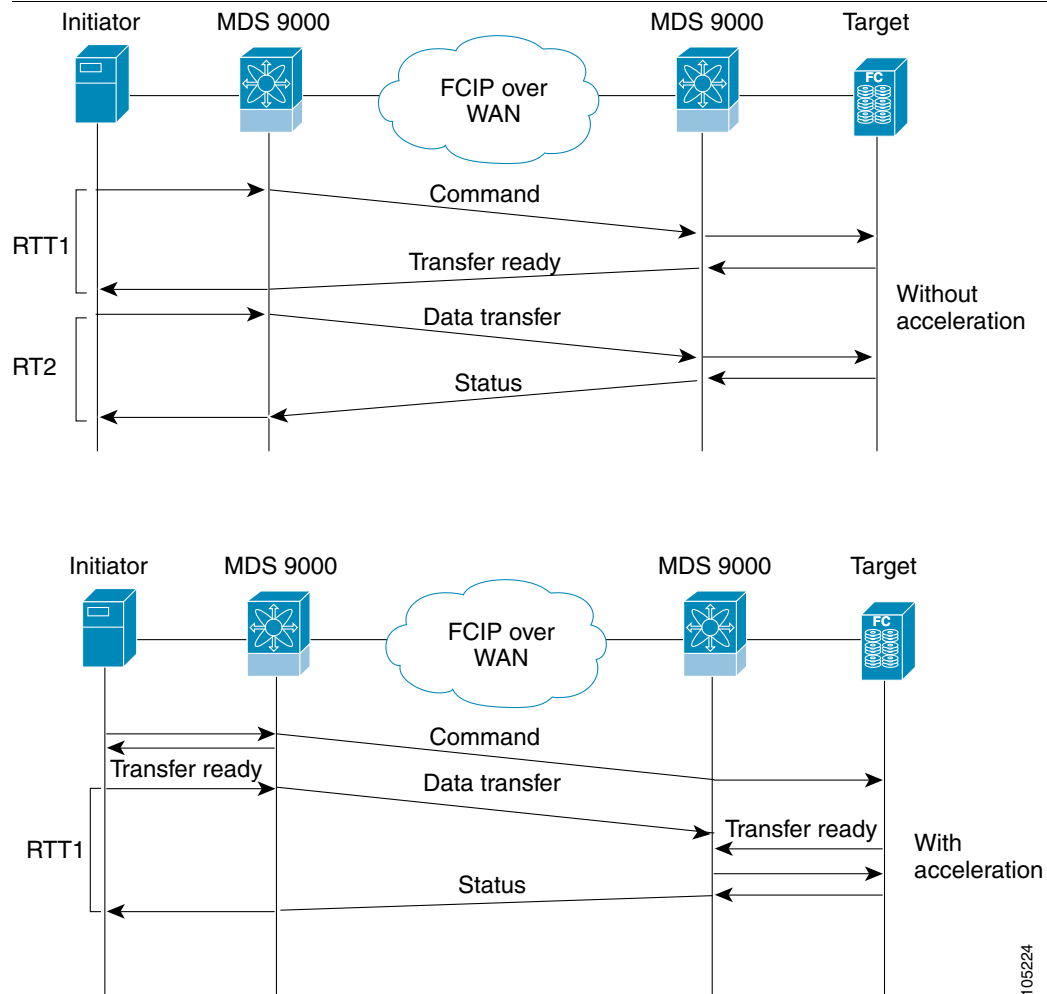
The WRITE command (see Figure 2-13), without write acceleration requires two round-trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Figure 2-13 FCIP Link Write Acceleration



Tip

Do not enable time stamp control on an FCIP interface with write acceleration configured.



105224

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

Write acceleration cannot be used across FSPF equal cost paths in FCIP deployments. Native Fibre Channel write acceleration can be used with PortChannels. Also, FCIP write acceleration can be used in PortChannels configured with channel mode active or constructed with PortChannel Protocol (PCP).

**Caution**

In Cisco MDS SAN-OS Release 2.0(1b) and later and NX-OS Release 4.x, FCIP write acceleration with FCIP ports as members of PortChannels are not compatible with the FCIP write acceleration in earlier releases.

FCIP Tape Acceleration

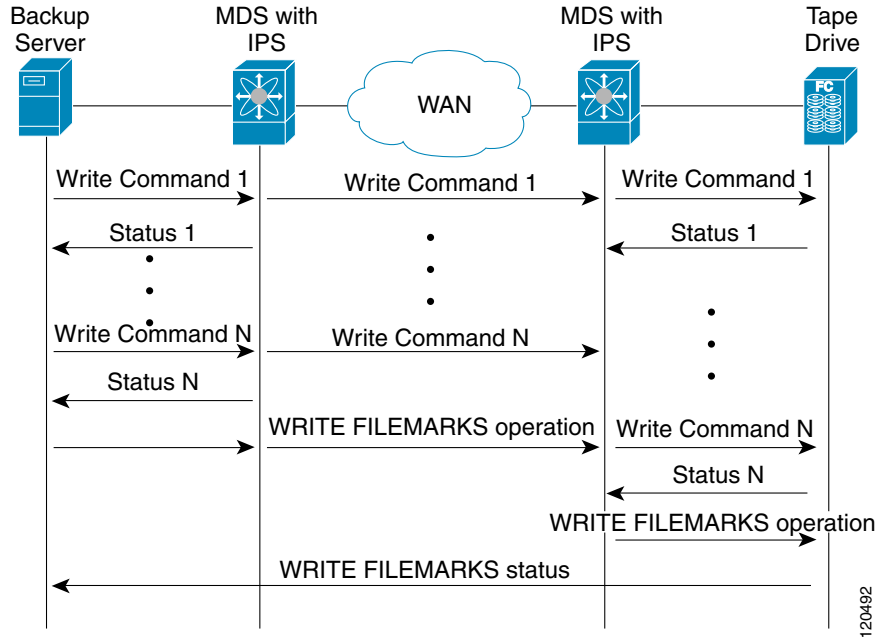
The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations. The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS NX-OS provides both tape write and read acceleration.

Applications that access tape drives normally have only one SCSI WRITE or READ operation outstanding to it. This single command process limits the benefit of the tape acceleration feature when using an FCIP tunnel over a long-distance WAN link. It impacts backup, restore, and restore performance because each SCSI WRITE or READ operation does not complete until the host receives a good status response from the tape drive. The FCIP tape acceleration feature helps solve this problem. It improves tape backup, archive, and restore operations by allowing faster data streaming between the host and tape drive over the WAN link.

In an example of tape acceleration for write operations, the backup server in (see [Figure 2-14](#)) issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the performance on WAN links.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-14 FCIP Link Tape Acceleration for Write Operations



At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.



Note


In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape library environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these sessions. You need to keep the FCIP link disabled for a couple of minutes before enabling the link. This does not apply to tape environments where the tape drives are either direct FC attached or exported as LUN 0.

The Cisco NX-OS provides reliable data delivery to the remote tape drives using TCP/IP over the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco NX-OS software.

In an example of tape acceleration for read operations, the restore server (see Figure 2-15) issues read operations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI read operations from the host, sends out SCSI read operations on its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI read operations from the host, sends out the cached data. This method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-15 FCIP Link Tape Acceleration for Read Operations



The Cisco NX-OS provides reliable data delivery to the restore application using TCP/IP over the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco NX-OS software recovers from any other errors.



Note

The FCIP tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tape acceleration feature is turned operationally off.



Tip

FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.



Caution

When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if an FCIP interface is up in a FICON VSAN, tape acceleration cannot be enabled on that interface.



Note

When you enable the tape acceleration feature for an FCIP tunnel, the tunnel is reinitialized and the write and read acceleration feature is also automatically enabled.

In tape acceleration for writes, after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch by not proxying the Transfer Ready. On completion of a write operation when some data buffers are freed, the local Cisco MDS switch resumes the proxying. Likewise, in tape acceleration for reads, after a certain

Send documentation comments to dcnm-san-docfeedback@cisco.com

amount of data has been buffered at the local Cisco MDS switch, the read operations to the tape drive are flow controlled by the remote Cisco MDS switch by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow control buffer size (the maximum buffer size is 12 MB).



Tip

We recommend that you use the default option for flow-control buffering.



Tip

Do not enable time-stamp control on an FCIP interface with tape acceleration configured.



Note

If one end of the FCIP tunnel is running Cisco MDS SAN-OS Release 3.0(1) or later and NX-OS Release 4.x, and the other end is running Cisco MDS SAN-OS Release 2.x, and tape acceleration is enabled, then the FCIP tunnel will run only tape write acceleration, not tape-read acceleration.



Note


In Cisco MDS NX-OS Release 4.2(1), the FCIP Tape Acceleration feature is not supported on FCIP back-to-back connectivity between MDS switches.

Tape Library LUN Mapping for FCIP Tape Acceleration

If a tape library provides logical unit (LU) mapping and FCIP tape acceleration is enabled, you must assign a unique LU number (LUN) to each physical tape drive accessible through a target port.

Figure 2-16 shows tape drives connected to Switch 2 through a single target port. If the tape library provides LUN mapping, then all the four tape drives should be assign unique LUNs.

Figure 2-16 FCIP LUN Mapping Example



Send documentation comments to dcnm-san-docfeedback@cisco.com

For the mappings described in [Table 2-1](#) and [Table 2-2](#), Host 1 has access to Drive 1 and Drive 2, and Host 2 has access to Drive 3 and Drive 4.

[Table 2-1](#) describes correct tape library LUN mapping.

Table 2-1 Correct LUN Mapping Example with Single Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 3	Drive 3
	LUN 4	Drive 4

[Table 2-2](#) describes incorrect tape library LUN mapping.

Table 2-2 Incorrect LUN Mapping Example with Single Hosts Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 1	Drive 3
	LUN 2	Drive 4

Another example setup is when a tape drive is shared by multiple hosts through a single tape port. For instance, Host 1 has access to Drive1 and Drive2, and Host 2 has access to Drive 2, Drive 3, and Drive 4. A correct LUN mapping configuration for such a setup is shown in [Table 2-3](#).

Table 2-3 Correct LUN Mapping Example with Multiple Host Access

Host	LUN Mapping	Drive
Host 1	LUN 1	Drive 1
	LUN 2	Drive 2
Host 2	LUN 2	Drive 2
	LUN 3	Drive 3
	LUN 4	Drive 4

FCIP Compression

The FCIP compression feature allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled. When enabled, the software defaults to using the **auto** mode (if a mode is not specified).



Note

The **auto** mode (default) selects the appropriate compression scheme based on the card type and bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 2-4 lists the modes used for different cards.

Table 2-4 Algorithm Classification

Mode	IPS Card	MPS 14/2 Card	MSM-18/4/MDS 9222i/SSN-16
mode1	SW	HW	HW
mode2	SW	SW	HW
mode3	SW	SW	HW



Note

With SAN-OS Release 3.3(1) and later and NX-OS Release 4.x, all compression options on the MDS 9222i switch and the MSM-18/4 module, mean hardware compression. Starting with Release 4.2(1), only auto compression and mode 2 compression are supported on the MDS 9222i switch, the MSM-18/4 module, and the SSN-16 module.

Table 2-5 lists the performance settings for different cards.

Table 2-5 Performance Settings

Bandwidth	IPS Cards	MPS 14/2 Card	MSM-18/4/MDS 9222i/SSN-16
Any	-	-	auto
>25 Mbps	mode 1	mode 1	auto
10-25 Mbps	mode 2	mode 2	auto
10 Mbps	mode 3	mode 3	auto



Note

The Cisco MDS 9216i and 9222i Switches also support the IP compression feature. The integrated supervisor module has the same hardware components that are available in the MPS-14/2 module.



Caution

The compression modes in Cisco SAN-OS Release 2.0(1b) and later and NX-OS Release 4.x are incompatible with the compression modes in Cisco SAN-OS Release 1.3(1) and earlier.



Tip

While upgrading from Cisco SAN-OS Release 1.x to Cisco SAN-OS Release 2.0(1b) or later and NX-OS Release 4.x, we recommend that you disable compression before the upgrade procedure, and then enable the required mode after the upgrade procedure.

If both ends of the FCIP link are running Cisco SAN-OS Release 2.0(1b) or later and NX-OS Release 4.x and you enable compression at one end of the FCIP tunnel, be sure to enable it at the other end of the link.

Default Settings

Table 2-6 lists the default settings for FCIP parameters.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 2-6 **Default FCIP Parameters**

Parameters	Default
TCP default port for FCIP	3225
minimum-retransmit-time	200 msec
Keepalive timeout	60 sec
Maximum retransmissions	4 retransmissions
PMTU discovery	Enabled
pmtu-enable reset-timeout	3600 sec
SACK	Enabled
max-bandwidth	1 Gbps
min-available-bandwidth	500 Mbps
round-trip-time	1 msec
Buffer size	0 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
TCP connection mode	Active mode is enabled
special-frame	Disabled
FCIP timestamp	Disabled
acceptable-diff range to accept packets	+/- 2000 msec
B port keepalive responses	Disabled
Write acceleration	Disabled
Tape acceleration	Disabled

Configuring FCIP

This section describes how to configure FCIP and includes the following topics:

- [Enabling FCIP, page 2-21](#)
- [Modifying an FCIP Link](#)
- [Creating FCIP Profiles, page 2-23](#)
- [Checking Trunk Status, page 2-24](#)
- [Launching Cisco Transport Controller, page 2-24](#)
- [Creating FCIP Links, page 2-25](#)
- [Configuring TCP Listener Ports, page 2-25](#)
- [Configuring TCP Parameters, page 2-26](#)
- [Assigning a Peer IP Address, page 2-31](#)
- [Configuring Active Connections, page 2-32](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Enabling Time Stamp Control](#), page 2-33
- [Configuring B Ports](#), page 2-34
- [Setting QoS Values](#), page 2-35
- [Configuring FCIP Write Acceleration](#), page 2-35
- [Configuring FCIP Tape Acceleration](#), page 2-36
- [Configuring FCIP Compression](#), page 2-37

Enabling FCIP

To begin configuring the FCIP feature, you must explicitly enable FCIP on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification operations commands for the FCIP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To use the FCIP feature, you need to obtain the SAN extension over IP package license (SAN_EXTN_OVER_IP or SAN_EXTN_OVER_IP_IPS4) (see the *Cisco Family NX-OS Licensing Guide*). By default, the MDS 9222i and 9216i switches are shipped with the SAN extension over IP package license.

Detailed Steps

To enable FCIP on any participating switch, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# feature fcip	Enables FCIP on that switch.
	switch(config)# no feature fcip	Disables (default) FCIP on that switch.



Note

In Cisco MDS SAN-OS Release 2.0 and later and NX-OS Release 4.x, there is an additional login prompt to log into a switch that is not a part of your existing fabric.

To create and manage FCIP links with DCNM-SAN, use the FCIP Wizard. Make sure that the IP Services Module is inserted in the required Cisco MDS 9000 Family switch, and that the Gigabit Ethernet interfaces on these switches are connected, and then verify the connectivity. The procedures for creating FCIP links using the FCIP Wizard are as follows:


- Select the endpoints.
- Choose the interfaces' IP addresses.
- Specify link attributes.
- (Optional) Enable FCIP write acceleration or FCIP compression.

To create FCIP links using the FCIP Wizard, follow these steps:

- Step 1** Click the **FCIP Wizard** icon in the DCNM-SAN toolbar (See Figure 2-17).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 2-17 FCIP Wizard



- Step 2** Choose the switches that act as endpoints for the FCIP link and click **Next**.
- Step 3** Choose the Gigabit Ethernet ports on each switch that will form the FCIP link.
- Step 4** If both Gigabit Ethernet ports are part of MPS-14/2 modules, check the **Enforce IPSEC Security** check box and set the **IKE Auth Key**. See the *Security Configuration Guide, Cisco DCNM for SAN* for information on IPsec and IKE.

Check the **Use Large MTU Size (Jumbo Frames)** option to use jumbo size frames of 2300. Since Fibre Channel frames are 2112, we recommended that you use this option. If you uncheck the box, the FCIP Wizard does not set the MTU size, and the default value of 1500 is set.



Note In Cisco MDS 9000 SAN-OS, Release 3.0(3), by default the **Use Large MTU Size (Jumbo Frames)** option is not selected.

- Step 5** Click **Next**.
You see the IP Address/Route input screen.
- Step 6** Select **Add IP Route** if you want to add an IP route, otherwise retain the defaults.
- Step 7** Click **Next**.
You see the TCP connection characteristics.
- Step 8** Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link.
You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.
- Step 9** Check the **Write Acceleration** check box to enable FCIP write acceleration on this FCIP link.
See the “FCIP Write Acceleration” section on page 2-12.
- Step 10** Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link.
See the “FCIP Compression” section on page 2-18.
- Step 11** Check the **Enable XRC Emulator** check box to enable XRC emulator on this FCIP link.
For more information on XRC Emulator, see the *Fabric Configuration Guide, Cisco DCNM for SAN*.
- Step 12** Click **Next**.
- Step 13** Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link.
- Step 14** Click **Finish** to create this FCIP link.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Modifying an FCIP Link

Once you have created FCIP links using the FCIP wizard, you may need to modify parameters for these links. This procedure includes modifying the FCIP profiles as well as the FCIP link parameters. Each Gigabit Ethernet interface can have three active FCIP links at one time.

To modify an FCIP link, follow these steps on both switches:

-
- Step 1** Configure the Gigabit Ethernet interface.
 - Step 2** Create an FCIP profile, and then assign the Gigabit Ethernet interface's IP address to the profile.
 - Step 3** Create an FCIP interface, and then assign the profile to the interface.
 - Step 4** Configure the peer IP address for the FCIP interface.
 - Step 5** Enable the interface.
-

Creating FCIP Profiles

Detailed Steps

To create an FCIP profile in switch 1 in Figure 2-5, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch1(config)# fcip profile 10 switch1(config-profile)#	Creates a profile for the FCIP connection. The valid range is from 1 to 255.
Step 3	switch1(config-profile)# ip address 10.100.1.25	Associates the profile (10) with the local IPv4 address of the Gigabit Ethernet interface (3/1).

To assign FCIP profile in switch 2 in Figure 2-5, follow these steps:

	Command	Purpose
Step 1	switch2# config terminal switch2(config)#	Enters configuration mode.
Step 2	switch2(config)# fcip profile 20 switch2(config-profile)#	Creates a profile for the FCIP connection.
Step 3	switch2(config-profile)# ip address 10.1.1.1	Associates the profile (20) with the local IPv4 address of the Gigabit Ethernet interface.

To create an FCIP profile in switch 1, follow these steps:

-
- Step 1** Verify that you are connected to a switch that contains an IPS module.
 - Step 2** From DCNM-SAN, choose **Switches > ISLs > FCIP** in the Physical Attributes pane. From Device Manager, choose **FCIP** from the IP menu.
 - Step 3** Click the **Create Row** button in DCNM-SAN or the **Create** button on Device Manager to add a new profile.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 4** Enter the profile ID in the ProfileId field.
 - Step 5** Enter the IP address of the interface to which you want to bind the profile.
 - Step 6** Modify the optional TCP parameters, if desired. Refer to DCNM for SAN Online Help for explanations of these fields.
 - Step 7** (Optional) Click the **Tunnels** tab and modify the remote IP address in the Remote IPAddress field for the endpoint to which you want to link.
 - Step 8** Enter the optional parameters, if required.
See the “FCIP Profiles” section on page 2-4 for information on displaying FCIP profile information.
 - Step 9** Click **Apply Changes** icon to save these changes.
-

Checking Trunk Status

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

Detailed Steps

To check the trunk status for the FCIP interface on Device Manager, follow these steps:

- Step 1** Make sure you are connected to a switch that contains an IPS module.
 - Step 2** Select **FCIP** from the IP menu.
 - Step 3** Click the **Trunk Config** tab if it is not already selected. You see the FCIP Trunk Config dialog box. This shows the status of the interface.
 - Step 4** Click the **Trunk Failures** tab if it is not already selected. You see the FCIP Trunk Failures dialog box.
-

Launching Cisco Transport Controller

Cisco Transport Controller (CTC) is a task-oriented tool used to install, provision, and maintain network elements. It is also used to troubleshoot and repair NE faults.

Detailed Steps

To launch CTC, follow these steps:

- Step 1** Right-click an ISL carrying optical traffic in the fabric.
- Step 2** Click **Element Manager**.
- Step 3** Enter the URL for the Cisco Transport Controller.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click **OK**.

Creating FCIP Links

Detailed Steps

To create an FCIP link endpoint in switch 1, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch(config)#	Enters configuration mode.
Step 2	switch1(config)# interface fcip 51 switch1(config-if)#	Creates an FCIP interface (51).
Step 3	switch1(config-if)# use-profile 10	Assigns the profile (10) to the FCIP interface.
Step 4	switch1(config-if)# peer-info ipaddr 10.1.1.1	Assigns the peer IPv4 address information (10.1.1.1 for switch 2) to the FCIP interface.
Step 5	switch1(config-if)# no shutdown	Enables the interface.

To create an FCIP link endpoint in switch 2, follow these steps:

	Command	Purpose
Step 1	switch2# config terminal switch2(config)#	Enters configuration mode.
Step 2	switch2(config)# interface fcip 52 switch2(config-if)#	Creates an FCIP interface (52).
Step 3	switch2(config-if)# use-profile 20	Binds the profile (20) to the FCIP interface.
Step 4	switch2(config-if)# peer-info ip address 10.100.1.25	Assigns the peer IPv4 address information (10.100.1.25 for switch 1) to the FCIP interface.
Step 5	switch1(config-if)# no shutdown	Enables the interface.

Configuring TCP Listener Ports

Detailed Steps

To configure TCP listener ports, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcip profile 20 switch(config-profile)#	Creates the profile (if it does not already exist) and enters profile configuration submenu. The valid range is from 1 to 255.

The default TCP port for FCIP is 3225. You can change this port by using the **port** command.

Send documentation comments to dcnm-san-docfeedback@cisco.com

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

You can control TCP behavior in a switch by configuring the TCP parameters that are described in this section.



Note

When FCIP is sent over a WAN link, the default TCP settings may not be appropriate. In such cases, we recommend that you tune the FCIP WAN link by modifying the TCP parameters (specifically bandwidth, round-trip times, and CWM burst size).

This section includes the following topics:

- [Configuring Minimum Retransmit Timeout, page 2-26](#)
- [Configuring Keepalive Timeout, page 2-26](#)
- [Configuring Maximum Retransmissions, page 2-27](#)
- [Configuring Path MTUs, page 2-27](#)
- [Configuring Selective Acknowledgments, page 2-28](#)
- [Configuring Window Management, page 2-28](#)
- [Configuring Monitoring Congestion, page 2-29](#)
- [Configuring Estimating Maximum Jitter, page 2-30](#)
- [Configuring Buffer Size, page 2-30](#)

Configuring Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds (msec).

To configure the minimum retransmit time, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp min-retransmit-time 500</code>	Specifies the minimum TCP retransmit time for the TCP connection to be 500 msec. The default is 200 msec and the range is from 200 to 5000 msec.
	<code>switch(config-profile)# no tcp min-retransmit-time 500</code>	Reverts the minimum TCP retransmit time to the factory default of 200 msec.

Configuring Keepalive Timeout

You can configure the interval that the TCP connection uses to verify that the FCIP link is functioning. This ensures that an FCIP link failure is detected quickly even when there is no traffic.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to tune the time taken to detect FCIP link failures.

You can configure the first interval during which the connection is idle (the default is 60 seconds). When the connection is idle for the configured interval, eight keepalive probes are sent at 1-second intervals. If no response is received for these eight probes and the connection remains idle throughout, that FCIP link is automatically closed.



Note

Only the first interval (during which the connection is idle) can be changed.

To configure the first keepalive timeout interval, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp keepalive-timeout 120</code>	Specifies the keepalive timeout interval for the TCP connection in seconds (120). The range is from 1 to 7200 seconds.
	<code>switch(config-profile)# no tcp keepalive-timeout 120</code>	Reverts the keepalive timeout interval to the default 60 seconds.

Configuring Maximum Retransmissions

You can specify the maximum number of times a packet is retransmitted before TCP decides to close the connection.

To configure maximum retransmissions, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp max-retransmissions 6</code>	Specifies the maximum number of retransmissions (6). 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.

Configuring Path MTUs

Path MTU (PMTU) is the minimum MTU on the IP network between the two endpoints of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a timeout of 3600 seconds. If TCP reduces the size of the maximum segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

Send documentation comments to dcnm-san-docfeedback@cisco.com

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 range is 60 to 3600 seconds.
	<code>switch(config-profile)# no tcp pmtu-enable reset-timeout 600</code>	Leaves PMTU discovery enabled but reverts the timeout to the default of 3600 seconds.

Configuring Selective Acknowledgments

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

To configure SACK, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp sack-enable</code>	Disables SACK.
	<code>switch(config-profile)# tcp sack-enable</code>	Enables SACK (default).

Configuring Window Management

The optimal TCP window size is automatically calculated using the maximum bandwidth parameter, the minimum available bandwidth parameter, and the dynamically measured round-trip time (RTT).



Note

The configured **round-trip-time** parameter determines the window scaling factor of the TCP connection. This parameter is only an approximation. The measured RTT value overrides the round trip time parameter for window management. If the configured **round-trip-time** is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

The **min-available-bandwidth** parameter and the measured RTT together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at minimum available bandwidth.

The **max-bandwidth-mbps** parameter and the measured RTT together determine the maximum window size.



Note

Set the maximum bandwidth to match the worst-case bandwidth available on the physical link, considering other traffic that might be going across this link (for example, other FCIP tunnels, WAN limitations). Maximum bandwidth should be the total bandwidth minus all other traffic going across that link.

Send documentation comments to dcnm-san-docfeedback@cisco.com

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 at 300 Mbps, and the RTT at 10 msec.
	<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 FCIP defaults are maximum bandwidth at 1 Gbps, minimum available bandwidth at 500 Mbps, and RTT at 1 msec.
	<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 available bandwidth at 2000 Kbps, and the RTT at 200 msec.

Configuring Monitoring Congestion

By enabling the congestion window monitoring (CWM) parameter, you allow TCP to monitor congestion after each idle period. The CWM parameter also determines the maximum burst size allowed after an idle period. By default, this parameter is enabled and the default burst size is 50 KB.

The interaction of bandwidth parameters and CWM and the resulting TCP behavior is outlined as follows:

- If the average rate of the Fibre Channel traffic over the preceding RTT is less than the min-available-bandwidth multiplied by the RTT, the entire burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.
- If the average rate of the Fibre Channel traffic is greater than min-available-bandwidth multiplied by the RTT, but less than max-bandwidth multiplied by the RTT, then if the Fibre Channel traffic is transmitted in burst sizes smaller than the configured CWM value the entire burst is sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the Fibre Channel traffic is larger than the min-available-bandwidth multiplied by the RTT and the burst size is greater than the CWM value, then only a part of the burst is sent immediately. The remainder is sent with the next RTT.

The software uses standard TCP rules to increase the window beyond the one required to maintain the min-available-bandwidth to reach the max-bandwidth.



Note

The default burst size is 50 KB.



Tip

We recommend that this feature remain enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To change the CWM defaults, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp cwm</code>	Disables congestion monitoring.
	<code>switch(config-profile)# tcp cwm</code>	Enables congestion monitoring and sets the burst size to its default size.
	<code>switch(config-profile)# tcp cwm burstsize 30</code>	Changes the burst size to 30 KB. The valid range is from 10 to 100 KB.
	<code>switch(config-profile)# no tcp cwm burstsize 25</code>	Leaves the CWM feature in an enabled state but changes the burst size to its factory default.

Configuring Estimating Maximum Jitter

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

You can configure the maximum estimated jitter in microseconds by the packet sender. The estimated variation should not include network queuing delay. By default, this parameter is enabled in Cisco MDS switches when IPS modules or MPS-14/2 modules are present.

The default value is 1000 microseconds for FCIP interfaces.

To configure the maximum jitter value, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp max-jitter</code>	Disables delay jitter estimation.
	<code>switch(config-profile)# tcp max-jitter</code>	Enables the delay jitter feature and sets the time to its factory default.
	<code>switch(config-profile)# tcp max-jitter 300</code>	Changes the time to 300 microseconds. The valid range is from 0 to 10000 microseconds.
	<code>switch(config-profile)# no tcp max-jitter 2500</code>	Leaves the delay jitter feature in an enabled state but changes the time to its factory default (1000 microseconds for FCIP interfaces).

Configuring Buffer Size

You can define the required additional buffering—beyond the normal send window size—that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default FCIP buffer size is 0 KB.



Note

Use the default if the FCIP traffic is passing through a high throughput WAN link. If you have a mismatch in speed between the Fibre Channel link and the WAN link, then time stamp errors occur in the DMA bridge. In such a situation, you can avoid time stamp errors by increasing the buffer size.

Send documentation comments to dcnm-san-docfeedback@cisco.com

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 16384 KB.
	<code>switch(config-profile)# no tcp send-buffer-size 5000</code>	Reverts the switch to its factory default. The default is 0 KB.

Assigning a 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. You can specify an IPv4 address or an IPv6 address.

Detailed Steps

To assign the peer information based on the IPv4 address and port number, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# peer-info ipaddr 10.1.1.1</code>	Assigns an IPv4 address to configure the peer information. Because no port is specified, the default port number (3225) is used.
	<code>switch(config-if)# no peer-info ipaddr 10.10.1.1</code>	Deletes the assigned peer port information.
Step 2	<code>switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000</code>	Assigns the IPv4 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.
	<code>switch(config-if)# no peer-info ipaddr 10.1.1.1 port 3000</code>	Deletes the assigned peer port information.
Step 3	<code>switch(config-if)# no shutdown</code>	Enables the interface.

To assign the peer information based on the IPv4 address and port number, follow these steps:

-
- Step 1** Expand ISLs and select **FCIP** in the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab. You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in DCNM-SAN or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the **NumTCPCon** field to the number of TCP connections from this FCIP link.
 - Step 8** (Optional) Check the **Enable** check box in the Time Stamp section and set the Tolerance field.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 9 (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.

To assign the peer information based on the IPv6 address and port number, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# peer-info ipaddr</code>	Assigns an IPv6 address to configure the peer information. Because no port is specified, the default port number (3225) is used.
	<code>switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a</code>	Deletes the assigned peer port information.
Step 2	<code>switch(config-if)# peer-info ipaddr 2001:0db8:800:200c::417a port 3000</code>	Assigns the IPv6 address and sets the peer TCP port to 3000. The valid port number range is 0 to 65535.
	<code>switch(config-if)# no peer-info ipaddr 2001:0db8:800:200c::417a port 3000</code>	Deletes the assigned peer port information.
Step 3	<code>switch(config-if)# ipv6 enable</code>	Enables IPv6 processing on the interface.
Step 4	<code>switch(config-if)# no shutdown</code>	Enables the interface.

To assign the peer information based on the IPv6 address and port number, follow these steps:

- Step 1** From DCNM-SAN, choose **ISLs > FCIP** from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**. You see the FCIP dialog box.
- Step 2** Click the **Tunnels** tab. You see the FCIP link information.
- Step 3** Click the **Create Row** icon in DCNM- SAN or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
- Step 4** Set the ProfileID and TunnelID fields.
- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
- Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
- Step 7** (Optional) Set the **NumTCPCon** field to the number of TCP connections from this FCIP link.
- Step 8** (Optional) Check the **Enable** check box in the Time Stamp section and set the Tolerance field.
- Step 9** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.

Configuring Active Connections

You can configure the required mode for initiating a TCP connection. By default, the active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection but waits for the peer to connect to it. By default, the switch tries two TCP connections for each FCIP link.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

To enable the passive mode, follow these steps:

	Command	Purpose
Step 1	switch(config-if)# passive-mode	Enables passive mode while attempting a TCP connection.
	switch(config-if)# no passive-mode	Reverts to the factory set default of using the active mode while attempting the TCP connection.
Step 2	switch(config-if)# no shutdown	Enables the interface.

Enabling Time Stamp Control

You can instruct the switch to discard packets that are outside the specified time. When enabled, this feature specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped.

By default, time stamp control is disabled in all switches in the Cisco MDS 9000 Family. If a packet arrives within a 2000 millisecond interval (+ or -2000 msec) from the network time, that packet is accepted.

**Note**

The default value for packet acceptance is 2000 microseconds. If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the *Cisco NX-OS Fundamentals Configuration Guide* for more information).

**Tip**

Do not enable time stamp control on an FCIP interface that has tape acceleration or write acceleration configured.

To enable or disable the time stamp control, follow these steps:

	Command	Purpose
Step 1	switch(config-if)# time-stamp Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link	Enables time stamp checking for received packets with a default acceptable time difference of 2000 msec.
	switch(config-if)# no time-stamp	Disables (default) time stamps.
Step 2	switch(config-if)# time-stamp acceptable-diff 4000	Configures the packet acceptance time. The valid range is from 500 to 10,000 msec.
	switch(config-if)# no time-stamp acceptable-diff 500	Deletes the configured time difference and reverts the difference to factory defaults. The default difference is a 2000-millisecond interval from the network time.
Step 3	switch(config-if)# no shutdown	Enables the interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring B Ports

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. B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface.

When an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

Detailed Steps

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.
	<code>switch(config-if)# no bport-keepalive</code>	Disables the reception of keepalive responses sent by a remote peer (default).

To enable B port mode, follow these steps:

-
- Step 1** Choose **ISLs > FCIP** from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**. You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab.
You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in DCNM-SAN or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the ProfileID and TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
 - Step 7** (Optional) Set the NumTCPCon field to the number of TCP connections from this FCIP link.
 - Step 8** Check the **Enable** check box in the B Port section of the dialog box and optionally check the **KeepAlive** check box if you want a response sent to an ELS Echo frame received from the FCIP peer.
 - Step 9** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Setting QoS Values

Detailed Steps

To set the QoS values, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# qos control 24 data 26</code>	Configures the control TCP connection and data connection to mark all packets on that DSCP value. The control and data value ranges from 0 to 63.
	<code>switch(config-if)# no qos control 24 data 26</code>	Reverts the switch to its factory default (marks all control and data packets with DSCP value 0).

Configuring FCIP Write Acceleration

To enable write acceleration, follow these steps:

	Command	Purpose
Step 1	<code>switch1# config terminal</code> <code>switch1(config)#</code>	Enters configuration mode.
Step 2	<code>switch1(config)# interface fcip 51</code> <code>switch1(config-if)#</code>	Creates an FCIP interface (51).
Step 3	<code>switch1(config-if)# write-accelerator</code> <code>switch1(config-if)# no write-accelerator</code>	Enables write acceleration. Disables write acceleration (default).

You can enable FCIP write acceleration when you create the FCIP link using the FCIP Wizard.

Detailed Steps

To enable write acceleration on an existing FCIP link, follow these steps:

-
- Step 1 Choose **ISLs > FCIP** from the Physical Attributes pane on DCNM-SAN.
You see the FCIP profiles and links in the Information pane.
On Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
 - Step 2 Click the **Tunnels (Advanced)** tab.
You see the FCIP link information.
 - Step 3 Check or uncheck the **Write Accelerator** check box.
 - Step 4 Choose the appropriate compression ratio from the **IP Compression** drop-down list.
 - Step 5 Click the **Apply Changes** icon to save these changes.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring FCIP Tape Acceleration

Detailed Steps

To enable FCIP tape acceleration, follow these steps:

	Command	Purpose
Step 1	<code>switch1# config terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch1(config)# interface fcip 5</code> <code>switch1(config-if)#</code>	Creates an FCIP interface (5).
Step 3	<code>switch1(config-if)# write-accelerator tape-accelerator</code>	Enables tape acceleration (and write acceleration—if not already enabled).
	<code>switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size auto</code>	Enables tape acceleration with automatic flow control (default).
	<code>switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size 2048</code>	Sets tape acceleration flow control buffer size to 2 MB.
	<code>switch1(config-if)# no write-accelerator tape-accelerator</code>	Disables tape acceleration (default) and resets the FCIP tunnel. Note The write acceleration feature remains enabled.
	<code>switch1(config-if)# no write-accelerator tape-accelerator flow-control-buffer-size 2048</code>	Changes the flow control buffer size to the default value of automatic. The tape acceleration and write acceleration features remain enabled. This command does not reset the FCIP tunnel.
	<code>switch1(config-if)# no write-accelerator</code>	Disables both the write acceleration and tape acceleration features and resets the FCIP tunnel.

To enable FCIP tape acceleration, follow these steps:

-
- Step 1** From DCNM-SAN, choose **ISLs > FCIP** from the Physical Attributes pane.
You see the FCIP profiles and links in the Information pane.
From Device Manager, choose **IP > FCIP**.
You see the FCIP dialog box.
 - Step 2** Click the **Tunnels** tab. You see the FCIP link information.
 - Step 3** Click the **Create Row** icon in DCNM-SAN or the **Create** button in Device Manager.
You see the FCIP Tunnels dialog box.
 - Step 4** Set the profile ID in the ProfileID field and the tunnel ID in the TunnelID fields.
 - Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
 - Step 6** Check the **TapeAccelerator** check box.
 - Step 7** (Optional) Set the other fields in this dialog box and click **Create** to create this FCIP link.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring FCIP Compression

Detailed Steps

To enable FCIP compression, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface fcip 51 switch(config-if)#	Creates an FCIP interface (51).
Step 3	switch(config-if)# ip-compression mode3	Enables high compression for low bandwidth links.
	switch(config-if)# ip-compression mode3	Defaults to using the auto mode.
	switch(config-if)# no ip-compression	Disables (default) the FCIP compression feature.

Verifying FCIP Configuration

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

Command	Purpose
show fcip profile 7	Displays FCIP profile configuration information.
show fcip summary	Displays FCIP summary.
show interface fcip 10	Displays the FCIP interface summary of counters for a specified interface.
show interface fcip 4 counters	Displays detailed FCIP interface standard counter information.
show interface fcip 51 description	Displays the FCIP interface description.
show interface fcip 3 counters brief	Displays brief FCIP interface counter information.
show fcip host-map 100	Displays exchanges processed by write acceleration at the specified host end FCIP link.
show fcip target-map 100	Displays exchanges processed by write acceleration at the specified target end FCIP link.
show interface fcip 4 counters	Displays detailed FCIP interface write acceleration counter information, if Enabled.
show fcip tape-session summary	Displays information about tapes for which exchanges are tape accelerated.
show fcip tape-session tunnel 1 host-end	Displays information about tapes for which exchanges are tape accelerated at the host-end FCIP link.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Command	Purpose
<code>show fcip tape-session tunnel 1 targ-end</code>	Displays information about tapes for which exchanges are tape accelerated at the target-end FCIP link.
<code>show interface fcip 1 counters</code>	Displays detailed FCIP interface tape acceleration counter information, if Enabled.
<code>show interface fcip 4 counters</code>	Displays detailed FCIP interface compression information, if Enabled
<code>show ips stats hw-comp all</code>	Displays the compression engine statistics for the MPS-14/2 module.

This section includes the following topics:

- Displaying FCIP Profile Configuration Information, page 2-39
- Displaying FCIP Profile Configuration Information, page 2-39
- Displaying FCIP Interface Information, page 2-39
- Displaying Write Acceleration Activity Information, page 2-41
- Displaying Tape Acceleration Activity Information, page 2-42
- Displaying FCIP Compression Information, page 2-44

Displaying FCIP Profile Information

Example 2-1 Displays FCIP Profiles

```
switch# show fcip profile
```

```
-----
ProfileId      Ipaddr         TcpPort
-----
1              10.10.100.150 3225
2              10.10.100.150 3226
40             40.1.1.2      3225
100            100.1.1.2     3225
200            200.1.1.2     3225
```

Example 2-2 Displays the Specified FCIP Profile Information

```
switch# show fcip profile 7
```

```
FCIP Profile 7
  Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
  Listen Port is 3225
  TCP parameters
    SACK is disabled
    PMTU discovery is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Send buffer size is 0 KB
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 15000 kbps
    Estimated round trip time is 1000 usec
```


Send documentation comments to dcnm-san-docfeedback@cisco.com

To verify the FCIP interfaces and Extended Link Protocol (ELP) on Device Manager, follow these steps:

-
- Step 1** Make sure you are connected to a switch that contains an IPS module.
 - Step 2** Select **FCIP** from the Interface menu.
 - Step 3** Click the **Interfaces** tab if it is not already selected. You see the FCIP Interfaces dialog box.
 - Step 4** Click the **ELP** tab if it is not already selected. You see the FCIP ELP dialog box.
-

Displaying FCIP Profile Configuration Information

Use the **show fcip profile** command to display FCIP profile configuration information.

```
switch# show fcip profile 7
FCIP Profile 7
  Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
  Listen Port is 3225
  TCP parameters
    SACK is disabled
    PMTU discovery is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Send buffer size is 0 KB
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 15000 kbps
    Estimated round trip time is 1000 usec
```

Displaying FCIP Profile Configuration Information

Use the **show fcip profile** command to display FCIP profile configuration information.

```
switch# show fcip profile 7
FCIP Profile 7
  Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
  Listen Port is 3225
  TCP parameters
    SACK is disabled
    PMTU discovery is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Send buffer size is 0 KB
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 15000 kbps
    Estimated round trip time is 1000 usec
```

Displaying FCIP Interface 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 Example 2-3 through Example 2-6.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 2-3 Displays the FCIP Summary

```
switch# show fcip summary
```

Tun	prof	Eth-if	peer-ip	Status	T	W	T	Enc	Comp	Bandwidth	rtt
					E	A	A			max/min	(us)
10	91	GE4/1	3.3.3.2	UP	N	N	N	N	N	1000M/1000M	2000
11	11	GE3/1.601	30.1.1.2	DOWN	N	N	N	N	N	1000M/500M	1000
12	12	GE3/1.602	30.1.2.2	DOWN	N	N	N	N	N	1000M/500M	1000
13	0		0.0.0.0	DOWN	N	N	N	N	N		
14	0		0.0.0.0	DOWN	N	N	N	N	N		
15	0		0.0.0.0	DOWN	N	N	N	N	N		
16	0		0.0.0.0	DOWN	N	N	N	N	N		
17	0		0.0.0.0	DOWN	N	N	N	N	N		
18	0		0.0.0.0	DOWN	N	N	N	N	N		
19	0		0.0.0.0	DOWN	N	N	N	N	N		
20	92	GE4/2	3.3.3.1	UP	N	N	N	N	N	1000M/1000M	2000
21	21	GE3/2.601	30.1.1.1	DOWN	N	N	N	N	N	1000M/500M	1000
22	22	GE3/2.602	30.1.2.1	DOWN	N	N	N	N	N	1000M/500M	1000

Example 2-4 Displays the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 10
fcip10 is up
  Hardware is GigabitEthernet
  Port WWN is 20:d0:00:0c:85:90:3e:80
  Peer port WWN is 20:d4:00:0c:85:90:3e:80
  Admin port mode is auto, trunk mode is on
  Port mode is E, FCID is 0x720000
  Port vsan is 91
  Speed is 1 Gbps
  Using Profile id 91 (interface GigabitEthernet4/1)
  Peer Information
    Peer Internet address is 3.3.3.2 and port is 3225
  Write acceleration mode is off
  Tape acceleration mode is off
  Tape Accelerator flow control buffer size is 256 KBytes
  IP Compression is disabled
  Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  QOS control code point is 0
  QOS data code point is 0
  B-port mode disabled
  TCP Connection Information
    50529025 Active TCP connections
      Local 0.0.0.7:6, Remote 0.0.0.200:0
    0 host table full 0 target entries in use
    211419104 Attempts for active connections, 1500 close of connections
  TCP Parameters
    Path MTU 124160 bytes
    Current retransmission timeout is 124160 ms
    Round trip time: Smoothed 127829 ms, Variance: 14336
    Advertized window: Current: 0 KB, Maximum: 14 KB, Scale: 14336
    Peer receive window: Current: 0 KB, Maximum: 0 KB, Scale: 51200
    Congestion window: Current: 14 KB, Slow start threshold: 49344 KB
    Current Send Buffer Size: 206463 KB, Requested Send Buffer Size: 429496728
  3 KB
    CWM Burst Size: 49344 KB
    5 minutes input rate 491913172779207224 bits/sec, 61489146597400903 bytes/sec, 0 frames/sec
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

5 minutes output rate 491913175298921320 bits/sec, 61489146912365165 bytes/s
ec, 14316551 frames/sec
  5702 frames input, 482288 bytes
    5697 Class F frames input, 481736 bytes
    5 Class 2/3 frames input, 552 bytes
    0 Reass frames
    0 Error frames timestamp error 0
  5704 frames output, 482868 bytes
    5698 Class F frames output, 482216 bytes
    6 Class 2/3 frames output, 652 bytes
    0 Error frames

```

Example 2-5 Displays Detailed FCIP Interface Standard Counter Information

```

switch# show interface fcip 4 counters
fcip4
  TCP Connection Information
...
  5 minutes input rate 207518944 bits/sec, 25939868 bytes/sec, 12471 frames/sec
  5 minutes output rate 205340328 bits/sec, 25667541 bytes/sec, 12340 frames/sec
  2239902537 frames input, 4658960377152 bytes
    18484 Class F frames input, 1558712 bytes
    2239884053 Class 2/3 frames input, 4658958818440 bytes
    0 Reass frames
    0 Error frames timestamp error 0
  2215051484 frames output, 4607270186816 bytes
    18484 Class F frames output, 1558616 bytes
    2215033000 Class 2/3 frames output, 4607268628200 bytes
    0 Error frames

```

Example 2-6 Displays the FCIP Interface Description

```

switch# show interface fcip 51 description
FCIP51
  Sample FCIP interface

```

The txbytes is the amount of data before compression. After compression, the compressed txbytes bytes are transmitted with compression and the uncompressed txbytes bytes are transmitted without compression. A packet may be transmitted without compression, if it becomes bigger after compression (see Example 2-7).

Example 2-7 Displays Brief FCIP Interface Counter Information

```

switch# show interface fcip 3 counters brief
-----
Interface           Input (rate is 5 min avg)      Output (rate is 5 min avg)
-----
                    Rate      Total                    Rate      Total
                    Mbits/s  Frames                    Mbits/s  Frames
-----
fcip3                9         0                        9         0
-----

```

Displaying Write Acceleration Activity Information

Example 2-8 through Example 2-10 show how to display information about write acceleration activity.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 2-8 Displays Exchanges Processed by Write Acceleration at the Specified Host End FCIP Link

```
switch# show fcip host-map 100

MAP TABLE (5 entries TOTAL entries 5)

  OXID | RXID | HOST FCID| TARG FCID| VSAN |  Index
-----+-----+-----+-----+-----+-----
0xd490|0xffff|0x00690400|0x00620426|0x0005|0x0000321f
0xd4a8|0xffff|0x00690400|0x00620426|0x0005|0x00003220
0xd4c0|0xffff|0x00690400|0x00620426|0x0005|0x00003221
0xd4d8|0xffff|0x00690400|0x00620426|0x0005|0x00003222
0xd4f0|0xffff|0x00690400|0x00620426|0x0005|0x00003223
```

Example 2-9 Displays Exchanges Processed by Write Acceleration at the Specified Target End FCIP Link

```
switch# show fcip target-map 100

MAP TABLE (3 entries TOTAL entries 3)

  OXID | RXID | HOST FCID| TARG FCID| VSAN |  Index
-----+-----+-----+-----+-----+-----
0xc308|0xffff|0x00690400|0x00620426|0x0005|0x00003364
0xc320|0xffff|0x00690400|0x00620426|0x0005|0x00003365
0xc338|0xffff|0x00690400|0x00620426|0x0005|0x00003366
```

Example 2-10 Displays Detailed FCIP Interface Write Acceleration Counter Information, if Enabled

```
switch# show interface fcip 4 counters
fcip4
  TCP Connection Information
...
  Write Accelerator statistics
    6091 packets in      5994 packets out
    0 frames dropped  0 CRC errors
    0 rejected due to table full
    0 ABTS sent      0 ABTS received
    0 tunnel synchronization errors
    37 writes recd      37 XFER_RDY sent (host)
    0 XFER_RDY rcvd (target)
    37 XFER_RDY rcvd (host)
    0 XFER_RDY not proxied due to flow control (host)
    0 bytes queued for sending
    0 estimated bytes queued on the other side for sending
    0 times TCP flow ctrl(target)
    0 bytes current TCP flow ctrl(target)
```

Displaying Tape Acceleration Activity Information

Example 2-11 through Example 2-14 show how to display information about tape acceleration activity.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 2-11 Displays Information About Tapes for Which Exchanges are Tape Accelerated

```
switch# show fcip tape-session summary
-----
Tunnel      Tunnel End    tape-fcid    lun          vsan         num-hosts
-----
1           host-end      EF0001       0x0002       0001         1
2           targ-end     650001       0x0003       0010         2
-----
```

Example 2-12 Displays Information About Tapes for Which Exchanges are Tape Accelerated at the Host-End FCIP Link

```
switch# show fcip tape-session tunnel 1 host-end

HOST TAPE SESSIONS (1 entries TOTAL entries 1)

Host Tape Session #1
  FCID 0xEF0001, VSAN 1, LUN 0x0002
  Outstanding Exchanges 0, Outstanding Writes 0
  Target End Write Buffering 0 Bytes, Auto Max Writes 3
  Flags 0x0, FSM state Non TA Mode
  Cached Reads 0
  First index 0xffffffff7, Last index 0xffffffff7, RA index 0x0000f99a
  Current index=0xffffffffe, Els Oxid 0xfff7
  Hosts 1
  FCID 0x770100
```

Example 2-13 Displays Information About Tapes for Which Exchanges are Tape Accelerated at the Target-End FCIP Link

```
switch# show fcip tape-session tunnel 1 targ-end

TARGET TAPE SESSIONS (1 entries TOTAL entries 1)

Target Tape Session #1
  FCID 0xEF0001, VSAN 1, LUN 0x0002
  Outstanding Exchanges 0, Outstanding Writes 0
  Host End Read Buffering 0 Bytes, Auto Max Read Blocks 3
  Flags 0x800, Timer Flags 0x0
  FSM State Default, Prev FSM State Bypass
  Relative Block offset 0
  First index 0xffffffff7, Last index 0xffffffff7, RA index 0x0000f99a
  Current index=0xffffffffe, Els Oxid 0xfff7
  Hosts 1
  FCID 0x770100
```

Example 2-14 Displays Detailed FCIP Interface Tape Acceleration Counter Information, if Enabled

```
switch# show interface fcip 1 counters
fcip1
  TCP Connection Information
  ....
  Tape Accelerator statistics
    1 Host Tape Sessions
    0 Target Tape Sessions
  Host End statistics
    Received 31521 writes, 31521 good status, 0 bad status
    Sent 31517 proxy status, 4 not proxied
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

Estimated Write buffer 0 writes 0 bytes
Received 31526 reads, 10 status
Sent 31516 cached reads
Read buffer 0 reads, 0 bytes
Host End error recovery statistics
Sent REC 0, received 0 ACCs, 0 Rejects
Sent ABTS 0, received 0 ACCs
Received 31 RECs, sent 2 ACCs, 0 Rejects
Received 0 SRRs, sent 0 ACCs, 0 Rejects
Received 0 TMF commands
Target End statistics
Received 0 writes, 0 good status, 0 bad status
Write Buffer 0 writes, 0 bytes
Received 0 reads, 0 good status, 0 bad status
Sent 0 reads, received 0 good status, 0 bad status
Sent 0 rewinds, received 0 good status, 0 bad status
Estimated Read buffer 0 reads, 0 bytes
Target End error recovery statistics
Sent REC 0, received 0 ACCs, 0 Rejects
Sent SRR 0, received 0 ACCs
Sent ABTS 0, received 0 ACCs
Received 0 TMF commands

```

Displaying FCIP Compression Information

Example 2-15 and Example 2-16 show how to display FCIP compression information.

Example 2-15 Displays Detailed FCIP Interface Compression Information, if Enabled

```

switch# show interface fcip 4 counters
fcip4
  TCP Connection Information
  ...
  IP compression statistics
    208752 rxbytes, 208752 rxbytes compressed
    5143584 txbytes
      0 txbytes compressed, 5143584 txbytes non-compressed
      1.00 tx compression ratio

```

Example 2-16 Displays the Compression Engine Statistics for the MPS-14/2 Module

```

switch# show ips stats hw-comp all
HW Compression Statistics for port GigabitEthernet3/1
  Compression stats
    0 input bytes, 0 output compressed bytes
    0 input pkts, 0 output compressed pkts
  Decompression stats
    0 input compressed bytes, 0 output bytes
    0 input compressed pkts, 0 output pkts
  Passthru stats
    0 input bytes, 0 output bytes
    0 input pkts, 0 output pkts
  Miscellaneous stats
    32 min input pktlen, 32 max input pktlen
    28 min output pktlen, 28 max output pktlen
    0 len mismatch, 0 incomplete processing
    0 invalid result, 0 invalid session drop
    0 comp expanded

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

HW Compression Statistics for port GigabitEthernet3/2
  Compression stats
    0 input bytes, 0 output compressed bytes
    0 input pkts, 0 output compressed pkts
  Decompression stats
    0 input compressed bytes, 0 output bytes
    0 input compressed pkts, 0 output pkts
  Passthru stats
    0 input bytes, 0 output bytes
    0 input pkts, 0 output pkts
  Miscellaneous stats
    32 min input pktlen, 32 max input pktlen
    28 min output pktlen, 28 max output pktlen
    0 len mismatch, 0 incomplete processing
    0 invalid result, 0 invalid session drop
    0 comp expanded

```

Field Descriptions for FCIP

This section describes the field description for FCIP.

FCIP Monitor

Field	Description
C3 Rx Bytes	The number of incoming bytes of data traffic.
C3 Tx Bytes	The number of outgoing bytes of data traffic.
CF Rx Bytes	The number of incoming bytes of control traffic.
CF Tx Bytes	The number of outgoing bytes of control traffic.
Rx Error	The number of inbound frames that contained errors preventing them from being deliverable to a higher-layer protocol.
Tx Error	The number of outbound frames that could not be transmitted because of errors.
RxDiscard	The number of inbound frames that were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol.
TxDiscard	The number of outbound frames that were chosen to be discarded even though no errors had been detected to prevent their being transmitted.

FCIP Interfaces Interfaces

Field	Description
Description	Alias name for the interface as specified by a network manager.
PortVsan	The VSAN ID to which this interface is statically assigned.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Oper	The current operating mode of the port.
AutoChannelCreate	If checked, automatically create the PortChannel.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
FailureCause	The cause of current operational state of the port.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value is N/A.
FICON Address	The FICON port address of this port.

FCIP Interfaces Trunk Failures

Field	Description
FailureCause	An entry is shown in this table if there is an error in the trunk status for the given VSAN.

FCIP FICON Configuration

Field	Description
Interface	This is a unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	This is the list of VSANs (in the range 1 through 2047) for which FICON tape acceleration is configured. Only VSANs with a cficonVsanEntry of CISCO-FICON-MIB present can be configured for FICON tape acceleration.
VSAN List Oper	This is the list of VSANs (in the range 1 through 2047) for which FICON tape acceleration is operationally ON.

FCIP Profiles

Field	Description
IP Address	The Internet address for this entity.
Port	A TCP port other than the FCIP well-known port on which the FCIP entity listens for new TCP connection requests.
SACK	Whether the TCP Selective Acknowledgement Option is enabled to allow the receiver end to acknowledge multiple lost frames in a single ACK, enabling faster recovery.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
KeepAlive (s)	The TCP keep-alive timeout for all links within this entity.
ReTrans MinTimeout (ms)	The TCP minimum retransmit timeout for all the links on this entity.
ReTrans Max	The maximum number of times that the same item of data will be retransmitted over a TCP connection. If delivery is not acknowledged after this number of retransmissions then the connection is terminated.
Send BufSize (KB)	The aggregate TCP send window for all TCP connections on all Links within this entity. This value is used for egress flow control. When the aggregate of the data queued on all connections within this entity reaches this value, the sender is flow controlled.
Bandwidth Max (Kb)	This is an estimate of the bandwidth of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
Bandwidth Min (Kb)	The minimum available bandwidth for the TCP connections on the links within this entity.
Est Round Trip Time (us)	This is an estimate of the round trip delay of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
PMTU Enable	The path MTU discovery.
PMTU ReserTimeout (sec)	The time interval for which the discovered path MTU is valid, before MSS reverts back to the negotiated TCP value.
CWM Enable	If true, congestion window monitoring is enabled.
CWM BurstSize (KB)	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.

FCIP Tunnels

Field	Description
Interface	This identifies the interface on this FCIP device to which this link pertains.
Attached	The interface on which this FCIP link was initiated.
B Port Enable	If true, the B port mode is enabled on the local FCIP link.
B Port KeepAlive	If true, a message is sent in response to a (Fibre Channel) ELS Echo frame received from the peer. Some B Port implementations use ELS Echo request/response frames as link keep alive.
Remote IP Address	The Internet address for the remote FCIP entity.
Remote TCP Port	The remote TCP port to which the local FCIP entity will connect if and when it initiates a TCP connection setup for this link.
Spc Frames Enable	If true, the TCP active opener initiates FCIP special frames and the TCP passive opener responds to the FCIP special frames. If it is set to false, the FCIP special frames are neither generated nor responded to.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Spc Frames RemoteWWN	The world wide name of the remote FC fabric entity. If this is a zero length string then this link would accept connections from any remote entity. If a WWN is specified then this link would accept connections from a remote entity with this WWN.
Spc Frames Remote Profile Id	The remote FCIP entity's identifier.

FCIP Tunnels (Advanced)

Field	Description
Interface	The interface on which this FCIP link was initiated.
Timestamp Enable	If true, the timestamp in FCIP header is to be checked.
Timestamp Tolerance	The accepted time difference between the local time and the timestamp value received in the FCIP header. By default this value will be EDTOV/2. EDTOV is the Error_Detect_Timeout Value used for Fibre channel Ports as the timeout value for detecting an error condition.
Number Connections	The maximum number of TCP connections allowed on this link.
Passive	If false, this link endpoint actively tries to connect to the peer. If true, the link endpoint waits for the peer to connect to it.
QoS Control	The value to be set for the ToS field in IP header for the TCP control connection.
QoS Data	The value to be set for the ToS field in IP header for the TCP Data connection.
IP Compression	What algorithm is used, if any.
Write Accelerator	The write accelerator allows for enhancing SCSI write performance.
Tape Accelerator	If true, the tape accelerator (which allows for enhancing Tape write performance) is enabled.
Tape Accelerator Oper	Write acceleration is enabled for the FCIP link.
TapeRead Accelerator Oper	Enabled automatically when the tape accelerator Oper is active.
FlowCtrlBufSize Tape (KB)	The size of the flow control buffer (64 K to 32 MB). If set to 0, flow control buffer size is calculated automatically by the switch.
IPSec	Indicates whether the IP security has been turned on or off on this link.
XRC Emulator	Check to enable XRC emulator. It is disabled by default.
XRC Emulator Oper	Indicates the operational status of XRC emulator.

Send documentation comments to dcnm-san-docfeedback@cisco.com

FCIP Tunnels (FICON TA)

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	The list of VSANs for which FICON tape acceleration is configured.
VSAN List Oper	The list of VSANs for which FICON tape acceleration is operationally on.

FCIP Tunnels Statistics

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
Rx IPCompRatio	The IP compression ratio for received packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.
Tx IPCompRatio	The IP compression ratio for transmitted packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.

FCIP XRC Statistics

Field	Description
ProfileId	Unique ID of the profile.
Interface	Name of the interface.
RRSAccelerated	The number of read record set IUs accelerated.
RRSForwarded	Number of read record set IUs forwarded.
BusyStatus	Number of instances of busy status received from the control unit.
UnitCheckStatus	Number of instances of unit check status received from the control unit.
cfmFcipLinkExtXRCEStats SelReset	Number of selective resets processed.
BufferAllocErrors	Number of buffer allocation errors.

Additional References

For additional information related to implementing FCIPs, see the following section:

- [Related Document, page 2-50](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Standards, page 2-50](#)
- [RFCs, page 2-50](#)
- [MIBs, page 2-50](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-FCIP-MGMT-EXT-MIB • CISCO-FCIP-MGMT-MIB • CISCO-FEATURE-CONTROL-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs

Feature History for FCIP

[Table 2-7](#) lists the release history for this feature. Only features that were introduced or modified in 5.0(1a) or a later release appear in the table.

Table 2-7 Feature History for FCIP

Feature Name	Releases	Feature Information
Configuring FCIP	5.0(1a)	FCIP Compression



Configuring the SAN Extension Tuner

The SAN Extension Tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. You can specify the size of the test I/O transfers and how many concurrent or serial I/Os to generate while testing. The SET reports the resulting I/Os per second (IOPS) and I/O latency, which helps you determine the number of concurrent I/Os needed to maximize FCIP throughput.

This chapter includes the following topics:

- [Information About the SAN Extension Tuner, page 3-1](#)
- [Licensing Requirements for SAN Extension Tuner, page 3-4](#)
- [Default Settings, page 3-4](#)
- [Configuring the SAN Extension Tuner, page 3-5](#)
- [Verifying SAN Extension Tuner Configuration, page 3-12](#)
- [Additional References, page 3-13](#)

Information About the SAN Extension Tuner

The SAN extension tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating either direct access (magnetic disk) or sequential access (magnetic tape) SCSI I/O commands and directing such traffic to a specific virtual target. Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options.



Note

SAN Extension Tuner is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem, the Cisco Fabric Switch for IBM BladeCenter, and 16-Port Storage Services Node (SSN-16).



Note

As of Cisco MDS SAN-OS Release 3.3(1a), SAN Extension Tuner is supported on the Multiservice Module (MSM) and the Multiservice Modular Switch.

Text Part Number:

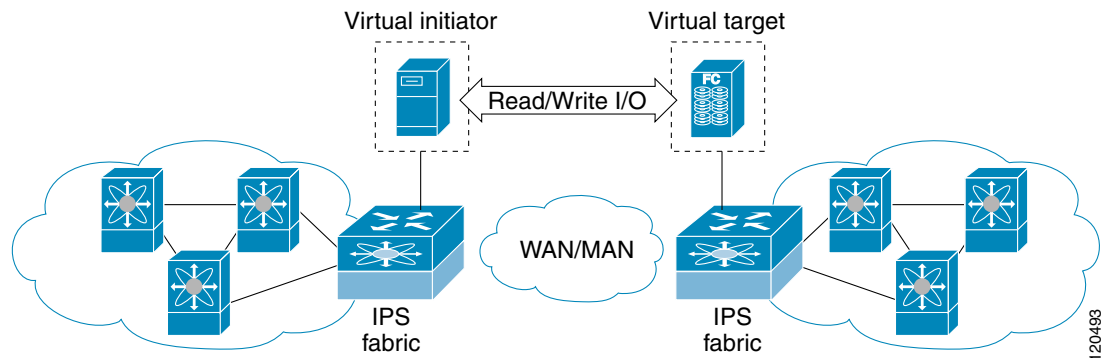
Send documentation comments to dcnm-san-docfeedback@cisco.com

Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. To achieve maximum throughput performance across the fabric, you can tune the following configuration parameters:

- The TCP parameters for the FCIP profile.
- The number of concurrent SCSI I/Os generated by the application.
- The transfer size used by the application over an FCIP link.

SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options (see Figure 3-1).

Figure 3-1 SCSI Command Generation to the Virtual Target



The SET feature assists with tuning by generating varying SCSI traffic workloads. It also measures throughput and response time per I/O over an FCIP link.

Before tuning the SAN fabric, be aware of the following guidelines:

- Following these implementation details:
 - The tuned configuration is not persistent.
 - The virtual N ports created do not register FC4 features supported with the name server. This is to avoid the hosts in the SAN from discovering these N ports as regular initiators or targets.
 - Login requests from other initiators in the SAN are rejected.
 - The virtual N ports do not implement the entire SCSI suite; it only implements the SCSI read and write commands.
 - Tuner initiators can only communicate with tuner targets.
- Verify that the Gigabit Ethernet interface is up at the physical layer (GBIC and Cable connected—an IP address is not required).
- Enable iSCSI on the switch (no other iSCSI configuration is required).
- **Enable the interface (no other iSCSI interface configuration is required)**
See “Creating iSCSI Interfaces” section on page 4-36 for more information.
- Create an iSCSI interface on the Gigabit Ethernet interface and enable the interface (no other iSCSI interface configuration is required)
see “Creating iSCSI Interfaces” section on page 4-36 for more information.
- Configure the virtual N ports in a separate VSAN or zone as required by your network.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Be aware that a separate VSAN with only virtual N ports is not required, but is recommended as some legacy HBAs may fail if logins to targets are rejected.
- Do not use same Gigabit Ethernet interface to configure virtual N ports and FCIP links—use different Gigabit Ethernet interfaces. While this is not a requirement, it is recommended as the traffic generated by the virtual N ports may interfere with the performance of the FCIP link.

This section includes the following topics:

- [SAN Extension Tuner Setup, page 3-3](#)
- [Data Pattern, page 3-4](#)

SAN Extension Tuner Setup

Figure 3-2 provides a sample physical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 3-2 N Port Tuning Configuration Physical Example





Figure 3-3 provides a sample logical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 3-3 Logical Example of N Port Tuning for a FCIP Link



[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Data Pattern

By default, an all-zero pattern is used as the pattern for data generated by the virtual N ports. You can optionally specify a file as the data pattern to be generated by selecting a data pattern file from one of three locations: the bootflash: directory, the volatile: directory, or the slot0: directory. This option is especially useful when testing compression over FCIP links. You can also use Canterbury corpus or artificial corpus files for benchmarking purposes.

Licensing Requirements for SAN Extension Tuner

To use the SET, you need to obtain the SAN_EXTN_OVER_IP license (see the *Cisco Family NX-OS Licensing Guide*).

The following table shows the licensing requirements for this feature:

License	License Description
SAN extension over IP package for IPS-8 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP_IPS4) 	It comprises the SAN extension tuner features.
SAN extension over IP package for MPS-14/2 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP_IPS2) 	This feature applies to the MPS-14/2 module and the fixed Cisco MDS 9216i Switch IP ports.
SAN extension over IP package for one MPS-18/4, one MPS-18/4 FIPS, or one SSN-16 engine in the Cisco MDS 9500 Series <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP_18_4) (SAN_EXTN_OVER_IP_SSN16) 	This feature applies to the MPS-18/4, MPS-18/4 FIPS, or SSN-16 modules.

Default Settings

Table 3-1 lists the default settings for tuning parameters.

Table 3-1 Default Tuning Parameters

Parameters	Default
Tuning	Disabled
Transfer ready size	Same as the transfer size in the SCSI write command
Outstanding I/Os	1
Number of transactions	1
Data generation format	All-zero format

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring the SAN Extension Tuner

This section includes the following topics:

- [Tuning the FCIP Link, page 3-5](#)
- [Using the SAN Extension Tuner Wizard, page 3-5](#)
- [Enabling the Tuner, page 3-7](#)
- [Configuring nWWN, page 3-7](#)
- [Configuring the Virtual N Port, page 3-7](#)
- [Assigning the SCSI Read/Write, page 3-8](#)
- [Assigning SCSI Tape Read/Write, page 3-10](#)
- [Configuring a Data Pattern, page 3-11](#)

Tuning the FCIP Link

Detailed Steps

To tune the required FCIP link, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Configure the nWWN for the virtual N ports on the switch. |
| Step 2 | Enable iSCSI on the interfaces on which you want to create the N ports. |
| Step 3 | Configure the virtual N ports on either side of the FCIP link. |
| Step 4 | Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see the <i>Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide</i>) to segregate the real initiators. Ensure that the zoning configuration is set up to allow the virtual N-ports to communicate with each other. |
| Step 5 | Start the SCSI read and write I/Os. |
| Step 6 | Add more N ports (as required) to other Gigabit Ethernet ports in the switch to obtain maximum throughput. One scenario that may require additional N ports is if you use FCIP PortChannels. |
-

Using the SAN Extension Tuner Wizard



You can use the SAN Extension Tuner wizard to perform the these tasks:

- [Configuring nWWN ports](#)
- [Enabling iSCSI](#)
- [Configuring Virtual N ports](#)
- [Assigning SCSI read and write CLI commands](#)
- [Assigning SCSI tape read and write CLI commands](#)
- [Configuring a data pattern for SCSI commands](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To tune the required FCIP link using the SAN Extension Tuner Wizard in Cisco DCNM-SAN, follow these steps:

-
- Step 1** Right-click a valid FCIP link in the Fabric pane, and then select **SAN Extension Tuner** from the drop-down list. You can also highlight the link and choose **Tools > Other > SAN Extension Tuner**.
You see the Select Ethernet Port Pair dialog box.
- Step 2** Select the Ethernet port pairs that correspond to the FCIP link you want to tune and click **Next**.
-  **Note** The Ethernet ports you select should be listed as down.
-
- You see the Specify Parameters dialog box.
- Step 3** Create and activate a new zone to ensure that the virtual N ports are not visible to real initiators in the SAN by clicking **Yes** to the zone creation dialog box.
- Step 4** (Optional) Change the default settings for the transfer data size and the number of concurrent SCSI read and write commands as follows:
- a. Set Transfer Size to the number of bytes that you expect your applications to use over the FCIP link.
 - b. Set Read I/O to the number of concurrent SCSI read commands you expect your applications to generate over the FCIP link.
 - c. Set Write I/O to the number of concurrent outstanding SCSI write commands you expect your applications to generate over the FCIP link.
-  **Note** There is only one outstanding I/O at a time to the virtual N port that emulates the tape behavior.
-
- d. Check the **Use Pattern File** check box and select a file that you want to use to set the data pattern that is generated by the SAN extension tuner. See the “Data Pattern” section on page 3-4.
- Step 5** Click **Next**.
You see the Results dialog box.
- Step 6** Click **Start** to start the tuner. The tuner sends a continuous stream of traffic until you click **Stop**.
- Step 7** Click **Show** to see the latest tuning statistics. You can select this while the tuner is running or after you stop it.
- Step 8** Click **Stop** to stop the SAN extension tuner.
-

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Enabling the Tuner

The tuning feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, tuning is globally enabled for the entire switch.

Detailed Steps

To enable the tuning feature, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# feature san-ext-tuner	Enables tuning.
	switch(config)# no feature san-ext-tuner	Removes the currently applied tuning configuration and disables tuning (default).

Configuring nWWN

Detailed Steps

To configure the nWWNs for the tuner in this switch, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submenu.
Step 2	switch(san-ext)# nWWN 10:00:00:00:00:00:00:00	Configures the nWWN for the SAN extension tuner.

Configuring the Virtual N Port

Detailed Steps

To configure the virtual N port for tuning, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# feature iscsi switch(config)# iscsi enable module 1	Enables iSCSI globally and then on module 1.
Step 3	switch(config)# interface iscsi 3/4 switch(config-if)#	Creates an iSCSI interface and enters interface configuration submenu.
Step 4	switch(config-if)# no shutdown	Enables the iSCSI interface.
Step 5	switch(config-if)# end switch#	Returns to EXEC mode.
Step 6	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submenu.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 7	<code>switch(san-ext)# nwwn 10:00:00:00:00:00:00:00</code>	Configures the nWWN for the SAN extension tuner.
Step 8	<code>switch(san-ext)# nport pwwn 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4</code> <code>switch(san-ext-nport)#</code>	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
	<code>switch(san-ext)# no nport pwwn 22:34:56:78:90:12:34:56 vsan 200 interface gigabitethernet 3/4</code>	Removes a virtual N port on the specified Gigabit Ethernet port and VSAN.

Assigning the SCSI Read/Write

You can assign SCSI read and write commands on a one-time basis or on a continuous basis.

Detailed Steps

To assign SCSI read or write commands on a one-time basis, follow these steps:

	Command	Purpose
Step 1	<code>switch# san-ext-tuner</code> <code>switch(san-ext)#</code>	Enters the SET configuration submode.
Step 2	<code>switch(san-ext)# nwwn 10:00:00:00:00:00:00:00</code>	Configures the nWWN for the SAN extension tuner.
Step 3	<code>switch(san-ext)# nport pwwn 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4</code> <code>switch(san-ext-nport)#</code>	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	<code>switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000</code>	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the read command. The total number of I/Os is 5,000,000 bytes.
Step 5	<code>switch(san-ext-nport)# write command-id 101 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000</code>	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the write command received by the target. The total number of I/Os is 5,000,000 bytes.
Step 6	<code>switch(san-ext-nport)# stop command-id 100</code>	Stops the command with the specified ID.
	<code>switch(san-ext-nport)# stop all</code>	Stops all outstanding commands.
Step 7	<code>switch(san-ext-nport)# clear counters</code>	Clears the counters associated with this N port.
Step 8	<code>switch(san-ext-nport)# end</code> <code>switch#</code>	Exits the SAN extension tuner submode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To generate SCSI read or write commands continuously, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submode.
Step 2	switch(san-ext)# nWWN 10:00:00:00:00:00:00:00	Configures the nWWN for the SAN extension tuner.
Step 3	switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)#	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous	Configures SCSI commands to be read continuously. Tip Use the stop command-id command to stop the outstanding configuration.
Step 5	switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous	Configures SCSI commands to be written continuously.
Step 6	switch(san-ext-nport)# stop command-id 100 switch(san-ext-nport)# stop command-id all	Stops the command with the specified ID. Stops all outstanding commands.
Step 7	switch(san-ext-nport)# clear counters	Clears the counters associated with this N port.
Step 8	switch(san-ext-nport)# end switch#	Exits the SAN extension tuner submode.

To specify a transfer ready size for a SCSI write command, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submode.
Step 2	switch(san-ext)# nWWN 10:00:00:00:00:00:00:00	Configures the nWWN for the SAN extension tuner.
Step 3	switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)#	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000	Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the write command received by the target. The total number of I/Os is 5,000,000 bytes.
Step 5	switch(san-ext-nport)# transfer-ready-size 512000 switch(san-ext-nport)# no transfer-ready-size 512000	Specifies the maximum transfer ready size of 512,000 bytes as a target for SCSI write commands. For a SCSI write command with a larger size, the target performs multiple transfers based on the specified transfer size. Removes the specified transfer ready size configuration for SCSI write commands.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 6	switch(san-ext-nport)# stop command-id 100	Stops the command with the specified ID.
Step 7	switch(san-ext-nport)# end switch#	Exits the SAN extension tuner submode.

Assigning SCSI Tape Read/Write

You can assign SCSI tape read and write commands on a one-time basis or on a continuous basis.



Note

There is only one outstanding I/O at a time to the virtual N-port that emulates the tape behavior.

Detailed Steps

To assign SCSI tape read and or write commands on a one-time basis, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submode.
Step 2	switch(san-ext)# nWWN 10:00:00:00:00:00:00:00	Configures the nWWN for the SAN extension tuner.
Step 3	switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)#	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	switch(san-ext-nport)# tape-read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 num-transactions 5000000 filemark-frequency 32	Specifies a transfer size of 512,000 bytes with space over the filemark every 32 SCSI read commands. The total number of I/Os is 5,000,000 bytes.
Step 5	switch(san-ext-nport)# tape-write command-id 101 target 22:22:22:22:22:22:22:22 transfer-size 512000 num-transactions 5000000 filemark-frequency 32	Specifies a transfer size of 512,000 bytes with filemarks written every 32 SCSI write commands. The total number of I/Os is 5,000,000 bytes.
Step 6	switch(san-ext-nport)# stop command-id 100	Stops the command with the specified ID.
	switch(san-ext-nport)# stop all	Stops all outstanding commands.
Step 7	switch(san-ext-nport)# clear counters	Clears the counters associated with this N port.
Step 8	switch(san-ext-nport)# end switch#	Exits the SAN extension tuner submode.

To generate SCSI tape read or write commands continuously, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submode.
Step 2	switch(san-ext)# nWWN 10:00:00:00:00:00:00:00	Configures the nWWN for the SAN extension tuner.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 3	switch(san-ext)# nport pWWN 12:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)#	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 4	switch(san-ext-nport)# tape-read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 continuous filemark-frequency 32	Configures SCSI tape read commands to be issued continuously. Tip Use the stop command-id command to stop the outstanding configuration.
Step 5	switch(san-ext-nport)# tape-write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 continuous filemark-frequency 32	Configures SCSI tape write commands to be issued continuously.
Step 6	switch(san-ext-nport)# stop command-id 100 switch(san-ext-nport)# stop command-id all	Stops the command with the specified ID. Stops all outstanding commands.
Step 7	switch(san-ext-nport)# clear counters	Clears the counters associated with this N port.
Step 8	switch(san-ext-nport)# end switch#	Exits the SAN extension tuner submode.

Configuring a Data Pattern

Detailed Steps

To optionally configure a data pattern for SCSI commands, follow these steps:

	Command	Purpose
Step 1	switch# san-ext-tuner switch(san-ext)#	Enters the SET configuration submode.
Step 2	switch(san-ext)# nport pWWN 12:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)#	Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as an initiator or a target.
Step 3	switch(san-ext-nport)# data-pattern-file bootflash://DataPatternFile	Specifies the data pattern sent by the virtual N port when it is a target for read commands and an initiator for write commands.. Tip This command should be configured on the target to change the data returned by read commands and on the initiator for write commands. This command is useful to define data sets which contain certain bit patterns or have certain compression ratios. The default data set of all zeros is very homogenous and very compressible.
	switch(san-ext-nport)# no data-pattern-file	Removes the specified data pattern configuration for SCSI read and write commands. The default is to send an all zero data pattern.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000	Specifies a transfer size of 512,000 bytes with two outstanding I/Os. The total number of I/Os is 5,000,000 bytes.
Step 5	switch(san-ext-nport)# stop command-id 100	Stops the command with the specified ID.
Step 6	switch(san-ext-nport)# clear counters	Clears the counters associated with this N port.
Step 7	switch(san-ext-nport)# end switch#	Exits the SAN extension tuner submode.

Verifying SAN Extension Tuner Configuration

To display SAN extension tuner configuration information, perform one of the following tasks:

Command	Purpose
<code>show flogi database</code>	Displays entries in the FLOGI database.
<code>show fcns database vsan 200</code>	Displays details for a VSAN entry in the FLOGI database.
<code>show san-ext-tuner interface gigabitethernet 3/4 nport pWWN 12:00:00:00:00:00:56 vsan 200 counters</code>	Displays all virtual N ports configured on the specified interface.
<code>show san-ext-tuner interface gigabitethernet 3/1</code>	Displays all N ports configured on the specified gigabit ethernet interface.
<code>show san-ext-tuner interface gigabitethernet 3/1 nport pWWN 10:0:0:0:0:0:1 vsan 91</code>	Displays the transfer ready size configured for a specified N port.
<code>show san-ext-tuner nports</code>	Displays all virtual N ports configured in this switch.

Verifying the SAN Extension Tuner Configuration

The `show` commands display the current SAN extension tuner settings for the Cisco MDS switch (see Examples 3-1 to 3-6).

Example 3-1 Displays Entries in the FLOGI Database

```
switch# show flogi database
-----
INTERFACE    VSAN    FCID      PORT NAME                               NODE NAME
-----
iscsi3/4     200     0x050000  12:00:00:00:00:00:00:56                 10:00:00:00:00:00:00:00
```

Example 3-2 Displays Details for a VSAN Entry in the FLOGI Database

```
switch# show fcns database vsan 200
VSAN 200
```


Send documentation comments to dcnm-san-docfeedback@cisco.com

```
-----
FCID          TYPE    PWWN (VENDOR)          FC4-TYPE:FEATURE
-----
0x020000      N       22:22:22:22:22:22:22  scsi-fcp
0x050000      N       12:00:00:00:00:00:56  scsi-fcp
-----
```

Example 3-3 *Displays All Virtual N Ports Configured on the Specified Interface*

```
switch# show san-ext-tuner interface gigabitethernet 3/4 nport pwwn
12:00:00:00:00:00:56 vsan 200 counters
Statistics for nport
Node name 10:00:00:00:00:00:00 Port name 12:00:00:00:00:00:56
I/Os per second          : 148
  Read                   : 0%
  Write                  : 100%
Ingress MB per second    : 0.02 MBs/sec (Max -0.02 MBs/sec)
Egress MB per second     : 73.97 MBs/sec (Max -75.47 MBs/sec)
Average Response time per I/O : Read - 0 us, Write - 13432 us
Maximum Response time per I/O : Read - 0 us, Write - 6953 us
Minimum Response time per I/O : Read - 0 us, Write - 19752 us
Errors                   : 0
```

Example 3-4 *Displays N Ports Configured on a Specified Gigabit Ethernet Interface*

```
switch# show san-ext-tuner interface gigabitethernet 3/1
-----
Interface          NODE NAME          PORT NAME          VSAN
-----
GigabitEthernet3/1 10:00:00:00:00:00:00 10:00:00:00:00:00:01 91
-----
```

Example 3-5 *Displays the Transfer Ready Size Configured for a Specified N Port*

```
switch# show san-ext-tuner interface gigabitethernet 3/1 nport pwwn 10:0:0:0:0:0:1 vsan
91
Node name          : 10:00:00:00:00:00:00
Port name          : 10:00:00:00:00:00:01
Transfer ready size : all
```

Example 3-6 *Displays All Virtual N Ports Configured in This Switch*

```
switch# show san-ext-tuner nports
-----
Interface          NODE NAME          PORT NAME          VSAN
-----
GigabitEthernet3/1 10:00:00:00:00:00:00 10:00:00:00:00:00:01 91
-----
```

Additional References

For additional information related to implementing FCIPs, see the following section:

- [Related Document, page 3-14](#)
- [Standards, page 3-14](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [RFCs, page 3-14](#)
- [MIBs, page 3-14](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



Configuring iSCSI

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note

The iSCSI feature is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

The Cisco MDS 9216i switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.



Note

For information on configuring Gigabit Ethernet interfaces, see [“Configuring Gigabit Ethernet Interface” section on page 7-5](#).

This chapter includes the following topics:

- [Information About iSCSI, page 4-2](#)
- [Licensing Requirements for iSCSI, page 4-32](#)
- [Guidelines and Limitations, page 4-32](#)
- [Default Settings, page 4-33](#)
- [Configuring iSCSI, page 4-34](#)
- [Configuring iSLB, page 4-56](#)
- [Distributing the iSLB Configuration Using CFS, page 4-67](#)
- [Configuring iSCSI Authentication, page 4-71](#)
- [Creating an iSNS Client Profile, page 4-75](#)
- [Configuring iSNS Cloud Discovery, page 4-80](#)
- [Verifying iSCSI Configuration, page 4-83](#)
- [Configuration Examples for iSCSI, page 4-105](#)
- [Field Descriptions for iSCSI, page 4-128](#)
- [Additional References, page 4-136](#)

Text Part Number:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Information About iSCSI

Cisco MDS 9000 Family IP Storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. To use the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric.

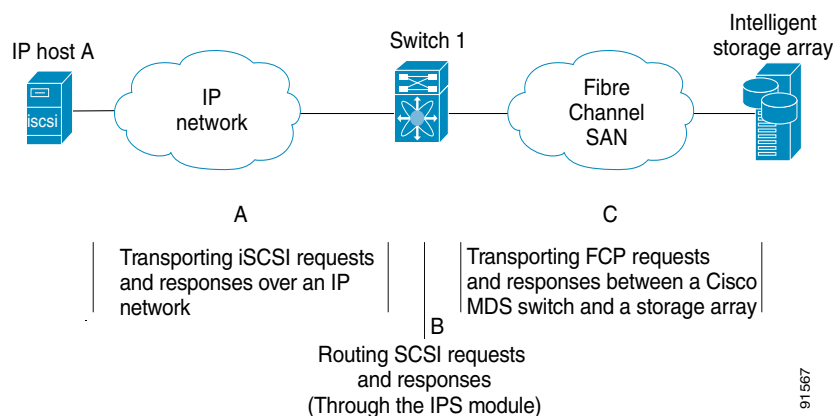


Note

The iSCSI feature is not supported on the Cisco Fabric Switch for HP c-Class Bladesystem and Cisco Fabric Switch for IBM BladeCenter.

The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see Figure 4-1).

Figure 4-1 Transporting iSCSI Requests and Responses for Transparent iSCSI Routing




Each iSCSI host that requires access to storage through the IPS module or MPS-14/2 module needs to have a compatible iSCSI driver installed. Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be an SCSI transport driver similar to a Fibre Channel driver in the host.

The IPS module or MPS-14/2 module provides transparent SCSI routing. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. It (see Figure 4-1) provides an example of a typical configuration of iSCSI hosts connected to an IPS module or MPS-14/2 module through the IP network access Fibre Channel storage on the Fibre Channel SAN.

The IPS module or MPS-14/2 module create a separate iSCSI SAN view and Fibre Channel SAN view. For the iSCSI SAN view, the IPS module or MPS-14/2 module creates iSCSI virtual targets and then maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical iSCSI targets were attached to the IP network (see Figure 4-2).


Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-2 *iSCSI SAN View—iSCSI Virtual Targets*




For the Fibre Channel SAN view, the IPS module or MPS-14/2 module presents iSCSI hosts as a virtual Fibre Channel host. The storage devices communicate with the virtual Fibre Channel host similar to communications performed with real Fibre Channel hosts (see Figure 4-3).

Figure 4-3 *Fibre Channel SAN View—iSCSI Host as an HBA*



The IPS modules or MPS-14/2 modules transparently map the command between the iSCSI virtual target and the virtual Fibre Channel host (see Figure 4-4).

Figure 4-4 *iSCSI to FCP (Fibre Channel) Routing*



Send documentation comments to dcnm-san-docfeedback@cisco.com

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions:

- The iSCSI requests and responses are transported over an IP network between the hosts and the IPS module or MPS-14/2 module.
- The SCSI requests and responses are routed between the hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). The IPS module or MPS-14/2 module performs this conversion and routing.
- The FCP requests or responses are transported between the IPS module or MPS-14/2 module and the Fibre Channel storage devices.

**Note**

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN. Refer to the IETF standards for IP storage at <http://www.ietf.org> for information on the iSCSI protocol.

About iSCSI Configuration Limits

iSCSI configuration has the following limits:

- The maximum number of iSCSI and iSLB initiators supported in a fabric is 2000.
- The maximum number of iSCSI and iSLB initiators supported is 200 per port.
- The maximum number of iSCSI and iSLB sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSCSI and iSLB session support by switch is 5000.
- The maximum number of iSCSI and iSLB targets supported in a fabric is 6000.

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module or MPS-14/2 module presents physical Fibre Channel targets as iSCSI virtual targets, allowing them to be accessed by iSCSI hosts. The module presents these targets in one of the two ways:

- **Dynamic mapping**—Automatically maps all the Fibre Channel target devices/ports as iSCSI devices. Use this mapping to create automatic iSCSI target names.
- **Static mapping**—Manually creates iSCSI target devices and maps them to the whole Fibre Channel target port or a subset of Fibre Channel LUNs. With this mapping, you must specify unique iSCSI target names.

Static mapping should be used when iSCSI hosts should be restricted to subsets of LUs in the Fibre Channel targets and/or iSCSI access control is needed (see the “[iSCSI Access Control](#)” section on [page 4-11](#)). Also, static mapping allows the configuration of transparent failover if the LUs of the Fibre Channel targets are reachable by redundant Fibre Channel ports (see the “[Transparent Target Failover](#)” section on [page 4-23](#)).

**Note**

The IPS module or MPS-14/2 module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module or MPS-14/2 module makes Fibre Channel targets available to iSCSI initiators.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Dynamic Mapping

When you configure dynamic mapping the IPS module or MPS-14/2 module imports all Fibre Channel targets to the iSCSI domain and maps each physical Fibre Channel target port as one iSCSI target. That is, all LUs accessible through the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the physical Fibre Channel target port.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

The IPS module or MPS-14/2 module creates an IQN formatted iSCSI target node name using the following conventions because the name must be unique in the SAN:

- IPS Gigabit Ethernet ports that are not part of a Virtual Router Redundancy Protocol (VRRP) group or PortChannel use this format:

```
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
```

- IPS ports that are part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
```

- Ports that are part of a PortChannel use this format:

```
iqn.1987-02.com.cisco:02.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
```




Note

If you have configured a switch name, then the switch name is used instead of the management IP address. If you have not configured a switch name, the management IP address is used.

With this convention, each IPS port in a Cisco MDS 9000 Family switch creates a unique iSCSI target node name for the same Fibre Channel target port in the SAN.

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0, LUN 1, and LUN 2, those LUNs would become available to an IP host through the iSCSI target node name `iqn.1987-05.com.cisco:05.MDS_switch_management_IP_address.01-01.3100112233445566` (see Figure 4-5).

Figure 4-5 Dynamic Target Mapping



Note

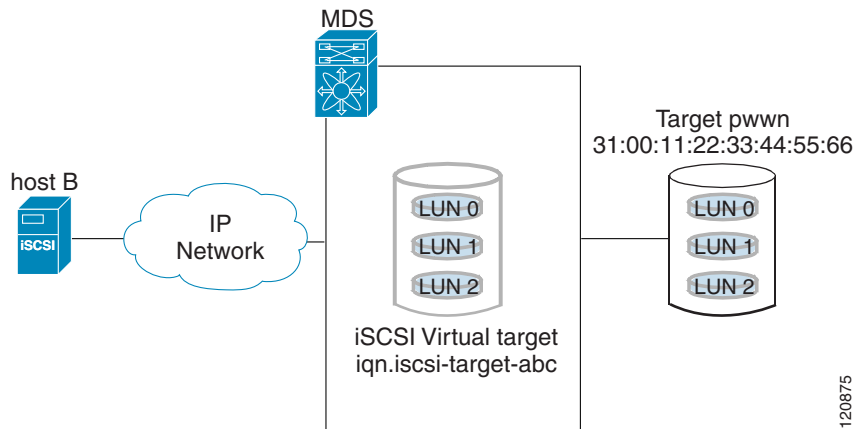
Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms (see the “iSCSI Access Control” section on page 4-11).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Static Mapping

You can manually (statically) create an iSCSI target by assigning a user-defined unique iSCSI node name to it. The iSCSI qualified name is restricted to a minimum length of 16 characters and a maximum of 223 characters. A statically mapped iSCSI target can either map the whole Fibre Channel target port (all LUNs in the target port mapped to the iSCSI target), or it can contain one or more LUs from a Fibre Channel target port (see Figure 4-6).

Figure 4-6 Statically Mapped iSCSI Targets



Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The IPS module or MPS-14/2 module connects to the Fibre Channel storage devices on behalf of the iSCSI host to send commands and transfer data to and from the storage devices. These modules use a virtual Fibre Channel N port to access the Fibre Channel storage devices on behalf of the iSCSI host. iSCSI hosts are identified by either iSCSI qualified name (IQN) or IP address.

Initiator Identification

iSCSI hosts can be identified by the IPS module or MPS-14/2 module using the following:

- iSCSI qualified name (IQN)

An iSCSI initiator is identified based on the iSCSI node name it provides in the iSCSI login. This mode can be useful if an iSCSI host has multiple IP addresses and you want to provide the same service independent of the IP address used by the host. An initiator with multiple IP addresses (multiple network interface cards—NICs) has one virtual N port on each IPS port to which it logs in.
- IP address

An iSCSI initiator is identified based on the IP address of the iSCSI host. This mode is useful if an iSCSI host has multiple IP addresses and you want to provide different service-based on the IP address used by the host. It is also easier to get the IP address of a host compared to getting the iSCSI node name. A virtual N port is created for each IP address it uses to log in to iSCSI targets. If the host using one IP address logs in to multiple IPS ports, each IPS port will create one virtual N port for that IP address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Initiator Presentation Modes

Two modes are available to present iSCSI hosts in the Fibre Channel fabric: transparent initiator mode and proxy initiator mode.

- In transparent initiator mode, each iSCSI host is presented as one virtual Fibre Channel host. The benefit of transparent mode is it allows a finer level of Fibre Channel access control configuration (similar to managing a “real” Fibre Channel host). Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.
- In proxy initiator mode, there is only one virtual Fibre Channel host per one IPS port and all iSCSI hosts use that to access Fibre Channel targets. In a scenario where the Fibre Channel storage device requires explicit LUN access control for every host, the static configuration for each iSCSI initiator can be overwhelming. In this case, using the proxy initiator mode simplifies the configuration.



Caution

Enabling proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [“Changing iSCSI Interface Parameters and the Impact on Load Balancing”](#) section on page 4-21.

The Cisco MDS switches support the following iSCSI session limits:

- The maximum number of iSCSI sessions on a switch is 5000.
- The maximum number of iSCSI sessions per IPS port in transparent initiator mode is 500.
- The maximum number of iSCSI sessions per IPS port in proxy initiator mode is 500.
- The maximum number of concurrent sessions an IPS port can create is five (but the total number of sessions that can be supported is 500).



Note

If more than five iSCSI sessions try to come up simultaneously on a port, the initiator receives a temporary error and later retries to create a session.


Transparent Initiator Mode

Each iSCSI host is presented as one virtual Fibre Channel host (that is, one Fibre Channel N port). The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration. Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.

When an iSCSI host connects to the IPS module or MPS-14/2 module, a virtual host N port (HBA port) is created for the host (see [Figure 4-7](#)). Every Fibre Channel N port requires a unique Node WWN and Port WWN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-7 Virtual Host HBA Port



After the virtual N port is created with the WWNs, a fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the virtual N port is online in the Fibre Channel SAN and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- IP address of the iSCSI host in the IP-address field on the name server
- IQN of the iSCSI host in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor-specific iSCSI GW flag in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server.

When all the iSCSI sessions from the iSCSI host are terminated, the IPS modules or MPS-14/2 modules perform an explicit Fabric logout (FLOGO) to remove the virtual N-port device from the Fibre Channel SAN (this indirectly de-registers the device from the Fibre Channel name server).

For every iSCSI session from the host to the iSCSI virtual target there is a corresponding Fibre Channel session to the real Fibre Channel target. There are three iSCSI hosts (see Figure 4-7), and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

WWN Assignment for iSCSI Initiators

An iSCSI host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

Send documentation comments to dcnm-san-docfeedback@cisco.com

Dynamic Mapping

With dynamic mapping, an iSCSI host is mapped to a dynamically generated port WWN (pWWN) and node WWN (nWWN). Each time the iSCSI host connects it might be mapped to a different WWN. Use this option if no access control is required on the Fibre Channel target device (because the target device access control is usually configured using the host WWN).

The WWNs are allocated from the MDS switch's WWN pool. The WWN mapping to the iSCSI host is maintained as long as the iSCSI host has at least one iSCSI session to the IPS port. When all iSCSI sessions from the host are terminated and the IPS module or MPS-14/2 module performs a FLOGO for the virtual N port of the host, the WWNs are released back to the switch's Fibre Channel WWN pool. These addresses are then available for assignment to other iSCSI hosts requiring access to the Fibre Channel Fabric.

The following are three dynamic initiator modes are supported:

- iSCSI—Dynamic initiators are treated as iSCSI initiators and can access dynamic virtual targets and configured iSCSI virtual targets.
- iSLB—Dynamic initiators are treated as iSLB initiators.
- Deny—Dynamic initiators are not allowed to log in to the MDS switch.

iSCSI dynamic mapping is the default mode of operation. This configuration is distributed using CFS.



Note

Configuring dynamic initiator modes is supported only through the CLI, not through Device Manager or Cisco DCNM for SAN.

Static Mapping

With static mapping, an iSCSI host is mapped to a specific pWWN and nWWN. This mapping is maintained in persistent storage and each time the iSCSI host connects, the same WWN mapping is used. This mode is required if you use access control on the target device.

You can implement static mapping in one of two ways:

- User assignment—You can specify your own unique WWN by providing them during the configuration process.
- System assignment—You can request that the switch provide a WWN from the switch's Fibre Channel WWN pool and keep the mapping in its configuration.



Tip

We recommend using the **system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the *Fabric Configuration Guide, Cisco DCNM for SAN* *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* for more information). You should not use any previously assigned WWNs.


Proxy Initiator Mode

In the event that the Fibre Channel storage device requires explicit LUN access control for every host use the transparent initiator mode (presenting one iSCSI host as one Fibre Channel host). Every iSCSI host has to be configured statically. This can mean several configuration tasks for each iSCSI host. If you do not need explicit LUN access control, using the proxy initiator mode simplifies the configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

In this mode, only one virtual host N port (HBA port) is created per IPS port. All the iSCSI hosts connecting to that IPS port will be multiplexed using the same virtual host N port (see Figure 4-8). This mode simplifies the task of statically binding WWNs. LUN mapping and assignment on the Fibre Channel storage array must be configured to allow access from the proxy virtual N port's pWWN for all LUNs used by each iSCSI initiator that connects through this IPS port. The LUN is then assigned to each iSCSI initiator by configuring iSCSI virtual targets (see the “Static Mapping” section on page 4-6) with LUN mapping and iSCSI access control (see the “iSCSI Access Control” section on page 4-11).

Figure 4-8 Multiplexing IPS Ports



Proxy initiator mode can be configured on a per IPS port basis, in which case only iSCSI initiators terminating on that IPS port will be in this mode.

When an IPS port is configured in proxy-initiator mode, fabric login (FLOGI) is done through the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the proxy-initiator virtual N port is online in the Fibre Channel fabric and virtual N port is registered in the Fibre Channel name server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel name server:

- iSCSI interface name iSCSI slot /port is registered in the symbolic-node-name field of the name server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor specific flag (iscsi-gw) in the FC-4 type field to identify the N-port device as an iSCSI gateway device in the name server

Similar to transparent initiator mode, the user can provide a pWWN and nWWN or request a system assigned WWN for the proxy initiator N port.



Caution

Enabling the proxy initiator mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the “Changing iSCSI Interface Parameters and the Impact on Load Balancing” section on page 4-21.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

VSAN Membership for iSCSI

VSAN membership can be configured for an iSCSI interface, called the port VSAN. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. The default port VSAN of an iSCSI interface is VSAN 1. Similar to Fibre Channel devices, iSCSI devices have two mechanisms by which VSAN membership can be defined.

- iSCSI host—VSAN membership to iSCSI host. (This method takes precedent over the iSCSI interface).
- iSCSI interface—VSAN membership to iSCSI interface. (All iSCSI hosts connecting to this iSCSI interface inherit the interface VSAN membership if the host is not configured in any VSAN by the iSCSI host method).

Advanced VSAN Membership for iSCSI Hosts

An iSCSI host can be a member of multiple VSANs. In this case, multiple virtual Fibre Channel hosts are created, one in each VSAN in which the iSCSI host is a member. This configuration is useful when certain resources such as Fibre Channel tape devices need to be shared among different VSANs.

iSCSI Access Control

Two methods of access control are available for iSCSI devices. Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel fabric, either or both of the access control methods can be used.

- Fiber Channel zoning-based access control—Fibre Channel zoning has been extended to support iSCSI devices, and this extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN. In the case of iSCSI, multiple iSCSI devices may be connected behind an iSCSI interface. Interface-based zoning may not be useful because all iSCSI devices behind the interface will automatically be within the same zone.
- iSCSI ACL-based access control—iSCSI-based access control is applicable only if static iSCSI virtual targets are created. For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets. By default, static iSCSI virtual targets are not accessible to any iSCSI host.

Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel fabric, either or both the access control mechanisms can be used.

The following topics are included in this section:

- [Fibre Channel Zoning-Based Access Control, page 4-11](#)
- [iSCSI-Based Access Control, page 4-12](#)
- [Enforcing Access Control, page 4-13](#)

Fibre Channel Zoning-Based Access Control

Cisco SAN-OS Release 3.x and NX-OS Release 4.1(1b) VSAN and zoning concepts have been extended to cover both Fibre Channel devices and iSCSI devices. Zoning is the standard access control mechanism for Fibre Channel devices, which is applied within the context of a VSAN. Fibre Channel zoning has been extended to support iSCSI devices, and this extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Common mechanisms for identifying members of a Fibre Channel zone are the following:

- Fibre Channel device pWWN.
- Interface and switch WWN. Device connecting via that interface is within the zone.

See the *Fabric Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS Fabric Configuration Guide for details on Fibre Channel zoning.

In the case of iSCSI, multiple iSCSI devices may be connected behind an iSCSI interface. Interface-based zoning may not be useful because all the iSCSI devices behind the interface will automatically be within the same zone.

In transparent initiator mode (where one Fibre Channel virtual N port is created for each iSCSI host as described in the “[Transparent Initiator Mode](#)” section on page 4-7), if an iSCSI host has static WWN mapping then the standard Fibre Channel device pWWN-based zoning membership mechanism can be used.

Zoning membership mechanism has been enhanced to add iSCSI devices to zones based on the following:

- IPv4 address/subnet mask
- IPv6 address/prefix length
- iSCSI qualified name (IQN)
- Symbolic-node-name (IQN)

For iSCSI hosts that do not have a static WWN mapping, the feature allows the IP address or iSCSI node name to be specified as zone members. Note that iSCSI hosts that have static WWN mapping can also use these features. IP address based zone membership allows multiple devices to be specified in one command by providing the subnet mask.



Note

In proxy initiator mode, all iSCSI devices connecting to an IPS port gain access to the Fibre Channel fabric through a single virtual Fibre Channel N port. Zoning based on the iSCSI node name or IP address will not have any effect. If zoning based on pWWN is used, then all iSCSI devices connecting to that IPS port will be put in the same zone. To implement individual initiator access control in proxy initiator mode, configure an iSCSI ACL on the virtual target (see the “[iSCSI-Based Access Control](#)” section on page 4-12).

iSCSI-Based Access Control

iSCSI-based access control is applicable only if static iSCSI virtual targets are created (see the “[Static Mapping](#)” section on page 4-6). For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets.

By default, static iSCSI virtual targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow an iSCSI virtual target to be accessed by all hosts. The initiator access list can contain one or more initiators. The iSCSI initiator can be identified by one of the following mechanisms:

- iSCSI node name
- IPv4 address and subnet
- IPv6 address

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

For a transparent mode iSCSI initiator, if both Fibre Channel zoning and iSCSI ACLs are used, then for every static iSCSI target that is accessible to the iSCSI host, the initiator's virtual N port should be in the same Fibre Channel zone as the Fibre Channel target.

Enforcing Access Control

IPS modules and MPS-14/2 modules use both iSCSI and Fibre Channel zoning-based access control lists to enforce access control. Access control is enforced both during the iSCSI discovery phase and the iSCSI session creation phase. Access control enforcement is not required during the I/O phase because the IPS module or MPS-14/2 module is responsible for the routing of iSCSI traffic to Fibre Channel.

- iSCSI discovery phase—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module or MPS-14/2 module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section. The IPS module or MPS-14/2 module does this by querying the Fibre Channel name server for all the devices in the same zone as the initiator in all VSANs. It then filters out the devices that are initiators by looking at the FC4-feature field of the FCNS entry. (If a device does not register as either initiator or target in the FC4-feature field, the IPS module or MPS-14/2 module will advertise it). It then responds to the iSCSI host with the list of targets. Each will have either a static iSCSI target name that you configure or a dynamic iSCSI target name that the IPS module or MPS-14/2 module creates for it (see the [“Dynamic Mapping” section on page 4-5](#)).
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module or MPS-14/2 module verifies if the specified iSCSI target (in the session login request) is allowed by both the access control mechanisms described in the [“iSCSI-Based Access Control” section on page 4-12](#).

If the iSCSI target is a static mapped target, the IPS module or MPS-14/2 module verifies if the iSCSI host is allowed within the access list of the iSCSI target. If the IP host does not have access, its login is rejected. If the iSCSI host is allowed, it validates if the virtual Fibre Channel N port used by the iSCSI host and the Fibre Channel target mapped to the static iSCSI virtual target are in the same Fibre Channel zone.

If the iSCSI target is an autogenerated iSCSI target, then the IPS module or MPS-14/2 module extracts the WWN of the Fibre Channel target from the iSCSI target name and verifies if the initiator and the Fibre Channel target is in the same Fibre Channel zone or not. If they are, then access is allowed.

The IPS module or MPS-14/2 module uses the Fibre Channel virtual N port of the iSCSI host and does a zone-enforced name server query for the Fibre Channel target WWN. If the FC ID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI Session Authentication

The IPS module or MPS-14/2 module supports the iSCSI authentication mechanism to authenticate the iSCSI hosts that request access to the storage devices. By default, the IPS modules or MPS-14/2 modules allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation, you can use any method supported and allowed by the Cisco MDS AAA infrastructure. AAA authentication supports a RADIUS, TACACS+, or local authentication device. See the *Security Configuration Guide, Cisco DCNM for SAN*.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

The **aaa authentication iscsi** command enables AAA authentication for the iSCSI host and specifies the method to use. See *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

iSCSI Immediate Data and Unsolicited Data Features

Cisco MDS switches support the iSCSI immediate data and unsolicited data features if requested by the initiator during the login negotiation phase. Immediate data is iSCSI write data contained in the data segment of an iSCSI command protocol data unit (PDU), such as combining the write command and write data together in one PDU. Unsolicited data is iSCSI write data that an initiator sends to the iSCSI target, such as an MDS switch, in an iSCSI data-out PDU without having to receive an explicit ready to transfer (R2T) PDU from the target.

These two features help reduce I/O time for small write commands because it removes one round-trip between the initiator and the target for the R2T PDU. As an iSCSI target, the MDS switch allows up to 64 KB of unsolicited data per command. This is controlled by the FirstBurstLength parameter during iSCSI login negotiation phase.

If an iSCSI initiator supports immediate data and unsolicited data features, these features are automatically enabled on the MDS switch with no configuration required.

Cisco MDS switches support the following advanced features for iSCSI interfaces:

- [iSCSI Listener Port, page 4-14](#)
- [TCP Tuning Parameters, page 4-14](#)
- [Setting QoS Values, page 4-55](#)
- [iSCSI Routing Modes, page 4-15](#)

iSCSI Listener Port

You can configure the TCP port number for the iSCSI interface that listens for new TCP connections. The default port number is 3260. Once you change the TCP port number, the iSCSI port only accepts TCP connections on the newly configured port.

TCP Tuning Parameters

You can configure the following TCP parameters:

- Minimum retransmit timeout (See the “Configuring Minimum Retransmit Timeout” section on page 2-26 for more information).
- Keepalive timeout.
- Maximum retransmissions (See the “Configuring Maximum Retransmissions” section on page 2-27 for more information).
- Path MTU (See the “Configuring Path MTUs” section on page 2-27 for more information).
- SACK (SACK is enabled by default for iSCSI TCP configurations).
- Window management (The iSCSI defaults are max-bandwidth is 1 Gbps, min-available-bandwidth is 70 Mbps, and round-trip-time is 1 msec). (See the “Configuring Window Management” section on page 2-28 for more information).
- Buffer size (The iSCSI default send buffer size is 4096 KB) (See the “Configuring Buffer Size” section on page 2-30 for more information).

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Window congestion monitoring (enabled by default and the default burst size is 50 KB) (See the “Configuring Monitoring Congestion” section on page 2-29 for more information).
- Maximum delay jitter (enabled by default and the default time is 500 microseconds).

iSCSI Routing Modes

Cisco MDS 9000 Family switches support multiple iSCSI routing modes. Each mode negotiates different operational parameters, has different advantages and disadvantages, and is suitable for different usages.

- Pass-thru mode

In pass-thru mode, the port on the IPS module or MPS 14/2 module converts and forwards read data frames from the Fibre Channel target to the iSCSI host frame-by-frame without buffering. This means that one data-in frame received is immediately sent out as one iSCSI data-in PDU.

In the opposite direction, the port on the IPS module or MPS 14/2 module limits the maximum size of iSCSI write data-out PDU that the iSCSI host can send to the maximum data size that the Fibre Channel target specifies that it can receive. The result is one iSCSI data-out PDU received sent out as one Fibre Channel data frame to the Fibre Channel target.

The absence of buffering in both directions leads to an advantage of lower forwarding latency. However, a small maximum data segment length usually results in lower data transfer performance from the host because of a higher processing overhead by the host system. Another benefit of this mode is iSCSI data digest can be enabled. This helps protect the integrity of iSCSI data carried in the PDU over what TCP checksum offers.

- Store-and-forward mode (default)

In store-and-forward mode, the port on the IPS module or MPS 14/2 module assembles all the Fibre Channel data frames of an exchange to build one large iSCSI data-in PDU before forwarding it to the iSCSI client.

In the opposite direction, the port on the IPS module or MPS 14/2 module does not impose a small data segment size on the host so the iSCSI host can send an iSCSI data-out PDU of any size (up to 256 KB). The port then waits until the whole iSCSI data-out PDU is received before it converts, or splits, the PDU, and forwards Fibre Channel frames to the Fibre Channel target.

The advantage of this mode is higher data transfer performance from the host. The disadvantages are higher transfer latency and that the iSCSI data digest (CRC) cannot be used.



Note The store-and-forward mode is the default forwarding mode.

- Cut-through mode

Cut-through mode improves the read operation performance over store-and-forward mode. The port on the IPS module or MPS 14/2 module achieves this by forwarding each Fibre Channel data-in frame to the iSCSI host as it is received without waiting for the whole exchange complete. There is no difference for write data-out operations from store-and-forward mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-9 compares the messages exchanged by the iSCSI routing modes.

Figure 4-9 *iSCSI Routing Modes*

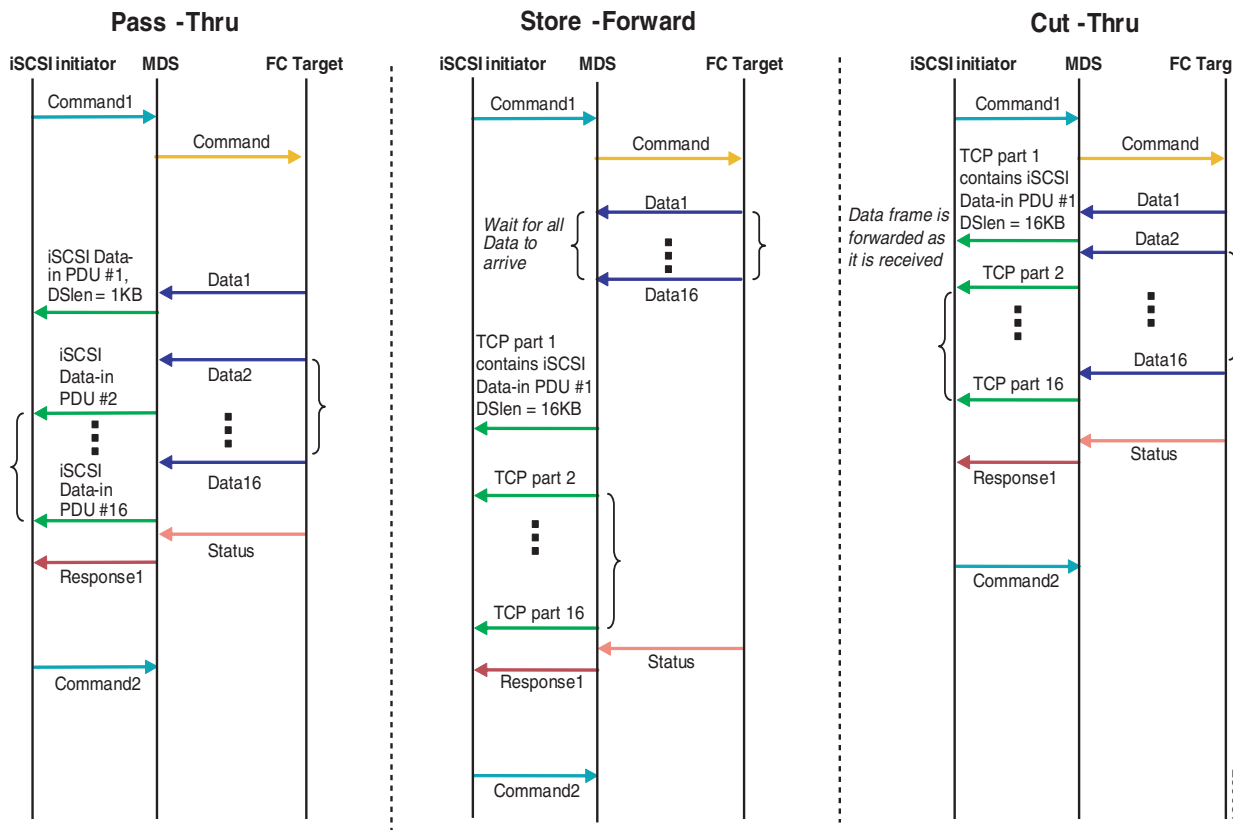


Table 4-1 compares the advantages and disadvantages of the different iSCSI routing modes.

Table 4-1 *Comparison of iSCSI Routing Modes*

Mode	Advantages	Disadvantages
Pass-thru	Low-latency Data digest can be used	Lower data transfer performance.
Store-and-forward	Higher data transfer performance	Data digest cannot be used.
Cut-thru	Improved read performance over store-and-forward	If the Fibre Channel target sent read data for different commands interchangeably, data of the first command is forwarded in cut-thru mode but the data of subsequent commands is buffered and the behavior is the same as store-and-forward mode. Data digest cannot be used.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Caution**

Changing the forwarding mode of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the “[Changing iSCSI Interface Parameters and the Impact on Load Balancing](#)” section on page 4-21.

About iSLB

The iSCSI server load balancing (iSLB) feature provides a means to easily configure large scale iSCSI deployments containing hundreds or even thousands of initiators. iSLB provides the following features:

- The iSLB initiator configuration is simplified with support for initiator targets and auto-zones.
- Cisco Fabric Services (CFS) eliminates the need for manual configuration by distributing the iSLB initiator configuration among all MDS switches in the fabric.
- Dynamic load balancing of iSLB initiators is available using iSCSI login redirect and VRRP.

When not using iSLB, configuring iSCSI requires the following:

- You need to perform multiple configuration steps on the MDS switch, including the following:
 - Initiator configuration using static pWWN and VSAN.
 - Zoning configuration for initiators and targets.
 - Optional create virtual target and give access to the initiator.
 - Configuration of target LUN mapping and masking on the storage system for the initiator based on the static pWWN created for the initiator on the MDS switch.
- You need to duplicate the configuration manually on multiple MDS switches.
- There is no load balancing for IPS ports. For example:
 - The Virtual Router Redundancy Protocol (VRRP) only supports active and backup, not load balancing.
 - You must use multiple VRRP groups and configure hosts in different groups.

iSLB provides the following features:

- The iSLB initiator configuration is simplified with support for initiator targets and auto-zones.
- Cisco Fabric Services (CFS) eliminates the need for manual configuration by distributing the iSLB initiator configuration among all MDS switches in the fabric.

**Note**

Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically mapped iSCSI initiator configurations are not distributed.

- Dynamic load balancing of iSLB initiators is available using iSCSI login redirect and VRRP.

About iSLB Initiators

iSLB initiators provide the following features in addition to those supported by iSCSI initiators:

- An iSLB initiator also supports iSLB virtual targets.
- Initiator targets—These targets are configured for a particular initiator.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Load balancing using iSCSI login redirect and VRRP—If iSCSI login redirect is enabled, the IPS Manager redirects incoming sessions to the best interface based on the calculated load for each interface.
- Configuration distribution to other switches using CFS.

iSLB initiators provide the following features in addition to those supported by iSCSI initiators:

- An iSLB initiator also supports iSLB virtual targets. These targets are very similar to iSCSI virtual targets with the exception that they do not include the advertise interface option and as a result are distributable using CFS.
- Initiator targets—These targets are configured for a particular initiator.
- Load balancing using iSCSI login redirect and VRRP—If load balancing is enabled, the IPS Manager redirects incoming sessions to the best interface based on the calculated load for each interface.
- Configuration distribution to other switches using CFS.

Assigning WWNs to iSLB Initiators

An iSLB host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping



Note

Assigning WWNs for iSLB initiators is the same as for iSCSI initiators. For information on dynamic and static mapping, see the [“WWN Assignment for iSCSI Initiators”](#) section on page 4-8.



Tip

We recommend using the **SystemAssign system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the *Fabric Configuration Guide, Cisco DCNM for SAN* *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* for more information). You should not use any previously assigned WWNs.

See the [“Configuring iSLB Using Device Manager”](#) procedure on page 4-56.

iSLB Initiator Targets

You can configure initiator targets using the device alias or the pWWN. You can also optionally specify one or more of the following optional parameters:

- Secondary pWWN
- Secondary device alias
- LUN mapping
- IQN
- VSAN identifier



Note

The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.

Send documentation comments to dcnm-san-docfeedback@cisco.com

In addition, you can disable auto-zoning.

If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

iSLB Session Authentication

The IPS module and MPS-14/2 module support the iSLB authentication mechanism to authenticate iSLB hosts that request access to storage. By default, the IPS module and MPS-14/2 module allow CHAP or None authentication of iSCSI initiators. If authentication is always used, you must configure the switch to allow only CHAP authentication.

For CHAP user name or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see the *Security Configuration Guide, Cisco DCNM for SAN* *Cisco MDS 9000 Family NX-OS Security Configuration Guide* for more information). AAA authentication supports RADIUS, TACACS+, or a local authentication device.



Note

Specifying the iSLB session authentication is the same as for iSCSI. See the “[iSCSI Session Authentication](#)” section on page 4-13.


About Load Balancing Using VRRP

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB. The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. The information is synchronized to all switches through CFS if recovery is needed when a master port fails. The initiator gets a temporary redirect iSCSI login response. The host then logs in to the backup port at its physical IP address. All iSCSI interfaces in a VRRP group that has load balancing enabled must have the same interface VSAN, authentication, proxy initiator mode, and forwarding mode.

You can configure Virtual Router Redundancy Protocol (VRRP) load balancing for iSLB. [Figure 4-10](#) shows an example of load balancing using iSLB.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-10 *iSLB Initiator Load Balancing Example*



The host is configured with a VRRP address as the portal address. When the VRRP master port receives the first iSCSI session from an initiator, it assigns a backup port to serve that particular host. This information is synchronized to all switches through CFS if recovery is needed when a master port fails. The initiator gets a temporary redirect iSCSI login response. The host then logs in to the backup port at its physical IP address. If the backup port goes down, the host will revert to the master port. The master port knows through CFS that the backup port has gone down and redirects the host to another backup port.



Note

If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.



Note

An initiator can also be redirected to the physical IP address of the master interface.



Tip

iSLB VRRP load balancing is based on the number of iSLB initiators and not number of sessions. Any iSLB initiator that has more targets configured than the other iSLB initiators (resulting in more sessions) should be configured with a higher load metric. For example, you can increase the load metric of the iSLB initiator with more targets to 3000 from the default value of 1000.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Caution**

A Gigabit Ethernet interface configured for iSLB can only be in one VRRP group because redirected sessions do not carry information about the VRRP IP address or group. This restriction allows the slave backup port to uniquely identify the VRRP group to which it belongs.

Changing iSCSI Interface Parameters and the Impact on Load Balancing

All iSCSI interfaces in a VRRP group that has load balancing enabled must have the same interface VSAN, authentication, proxy initiator mode, and forwarding mode. When you need to change any of these parameters for the iSCSI interfaces in a VRRP group, you must do so one interface at a time. During the transition time when the parameter is changed on some interfaces in the VRRP group and not the others, the master port does not redirect new initiators and instead handles them locally.

**Caution**

Changing the VSAN, proxy initiator, authentication, and forwarding mode for iSCSI interfaces in a VRRP group can cause sessions to go down multiple times.

VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces

When the VRRP master receives an iSCSI session request from an initiator, it first checks for an existing mapping to one of the interfaces in that VRRP group. If such a mapping exists, the VRRP master redirects the initiator to that interface. If no such mapping exists, the VRRP master selects the least loaded interface and updates the selected interface's load with the initiator's iSLB metric (weight).

**Note**

The VRRP master interface is treated specially and it needs to take a lower load compared to the other interfaces. This is to account for the redirection work performed by the master interface for every session. A new initiator is assigned to the master interface only if the following is true for every other interface:

$$\text{VRRP backup interface load} > [2 * \text{VRRP master interface load} + 1]$$

About iSLB Configuration Distribution Using CFS

You can distribute the configuration for iSLB initiators and initiator targets on an MDS switch. This feature lets you synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, global authentication, and iSCSI dynamic initiator mode parameters are also distributed. CFS distribution is disabled by default.

Configuration for iSLB initiators and initiator targets on an MDS switch can be distributed using the Cisco Fabric Services (CFS). This feature allows you to synchronize the iSLB configuration across the fabric from the console of a single MDS switch. The iSCSI initiator idle timeout, iSCSI dynamic initiator mode, and global authentication parameters are also distributed. CFS distribution is disabled by default (see the *System Management Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS System Management Configuration Guide for more information).

After enabling the distribution, the first configuration starts an implicit session. All server configuration changes entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database.

Send documentation comments to dcnm-san-docfeedback@cisco.com

When CFS is enabled for iSLB, the first iSLB configuration operation starts a CFS session and locks the iSLB configuration in the fabric. The configuration changes are applied to the pending configuration database. When you make the changes to the fabric, the pending configuration is distributed to all the switches in the fabric. Each switch then validates the configuration. This check ensures the following:

- The VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names do not conflict with the iSCSI initiators on all the switches.

After the check completes successfully, all the switches commit the pending configuration to the running configuration. If any check fails, the entire commit fails.



Note

iSLB is only fully supported when CFS is enabled. Using iSLB auto-zoning without enabling CFS mode may cause traffic disruption when any zone set is activated.



Note

CFS does not distribute non-iSLB initiator configurations or import Fibre Channel target settings.

Non-iSLB virtual targets will continue to support advertised interfaces option.



Tip

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

Locking the Fabric

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

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Modifications from this point on are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.



Note

iSCSI configuration changes are not allowed when an iSLB CFS session is active.

CFS Merge Process

When two fabrics merge, CFS attempts to merge the iSLB configuration from both the fabrics. A designated switch (called the *dominant switch*) in one fabric sends its iSLB configuration to a designated switch (called the *subordinate switch*) in the other fabric. The subordinate switch compares its running configuration to the received configuration for any conflicts. If no conflicts are detected, it merges the two configurations and sends it to all the switches in both the fabrics. Each switch then validates the configuration. This check ensures the following:

- VSANs assigned to the iSLB initiators are configured on all the switches.
- The static WWNs configured for the iSLB initiators are unique and available on all the switches.
- The iSLB initiator node names have no conflicts with iSCSI initiators on all the switches.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

If this check completes successfully, the subordinate switch directs all the switches to commit the merged configuration to running configuration. If any check fails, the merge fails.

The **show islb merge status** command displays the exact reason for the failure. The first successful commit request after a merge failure takes the fabric out of the merge failure state.

iSLB CFS Merge Status Conflicts

Merge conflicts may occur. User intervention is required for the following merge conflicts:

- The iSCSI global authentication or iSCSI initiator idle timeout parameters are not configured the same in the two fabrics.
- The same iSLB initiator is configured differently in the two fabrics.
- An iSLB initiator in one fabric has the same name as an iSCSI initiator in the other fabric.
- Duplicate pWWN/nWWN configuration is detected in the two fabric. For example, a pWWN/nWWN configured for an iSLB initiator on one fabric is configured for an iSCSI initiator or a different iSLB initiator in the other fabric.
- A VSAN configured for an iSLB initiator in one fabric does not exist in the other fabric.



Tip

Check the syslog for details on merge conflicts.

User intervention is not required when the same iSLB initiator has a different set of non-conflicting initiator targets. The merged configuration is the union of all the initiator targets.

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- [Transparent Target Failover, page 4-23](#)
- [iSCSI High Availability with Host Running Multi-Path Software, page 4-23](#)
- [iSCSI HA with Host Not Having Any Multi-Path Software, page 4-24](#)
- [LUN Trespass for Storage Port Failover, page 4-25](#)

Transparent Target Failover

The following high availability features are available for iSCSI configurations:

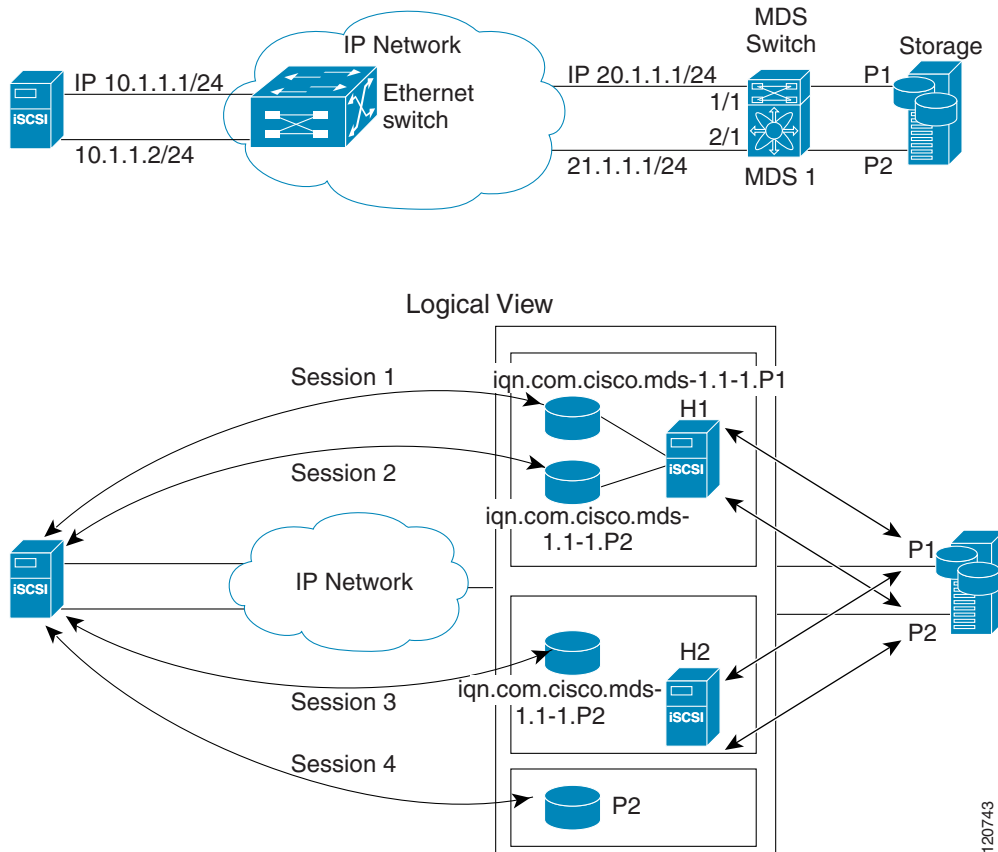
- iSCSI high availability with host running multi-path software—In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load balancing or failover across the different paths to access the storage.
- iSCSI high availability with host not having multi-path software—Without multi-path software, the host does not have knowledge of the multiple paths to the same storage.

iSCSI High Availability with Host Running Multi-Path Software

Figure 4-11 shows the physical and logical topology for an iSCSI HA solution for hosts running multi-path software. In this scenario, the host has four iSCSI sessions. There are two iSCSI sessions from each host NIC to the two IPS ports.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-11 Host Running Multi-Path Software



Each IPS ports is exporting the same two Fibre Channel target ports of the storage but as different iSCSI target names if you use dynamic iSCSI targets). So the two IPS ports are exporting a total of four iSCSI target devices. These four iSCSI targets map the same two ports of the Fibre Channel target.

The iSCSI host uses NIC-1 to connect to IPS port 1 and NIC-2 to connect to IPS port 2. Each IPS port exports two iSCSI targets, so the iSCSI host creates four iSCSI sessions.

If the iSCSI host NIC-1 fails (see Figure 4-11 for the physical view), then sessions 1 and 2 fail but we still have sessions 3 and 4.

If the IPS port 1 fails, the iSCSI host cannot connect to the IPS port, and sessions 1 and 2 fail. But sessions 3 and 4 are still available.

If the storage port 1 fails, then the IPS ports will terminate sessions 1 and 3 (put iSCSI virtual target `iqn.com.cisco.mds-5.1-2.p1` and `iqn-com.cisco.mds-5.1-1.p1` in offline state). But sessions 2 and 4 are still available.

In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load-balancing or failover across the different paths to access the storage.

iSCSI HA with Host Not Having Any Multi-Path Software

The above topology will not work if the host does not have multi-path software because the host has multiple sessions to the same storage. Without multi-path software the host does not have knowledge of the multiple paths to the same storage.


Send documentation comments to dcnm-san-docfeedback@cisco.com

IP storage has two additional features that provide an HA solution in this scenario.

- IPS ports support the VRRP feature
- IPS has transparent Fibre Channel target failover for iSCSI static virtual targets.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see Figure 4-12).

Figure 4-12 Static Target Importing Through Two Fibre Channel Ports



In Figure 4-12, you can create an iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

The failover to a secondary port is done transparently by the IPS port without impacting the iSCSI session from the host. All outstanding I/Os are terminated with a check condition status when the primary port fails. New I/Os received during the failover are not completed and receive a busy status.



Tip

If you use LUN mapping, you can define a different secondary Fibre Channel LUN if the LU number is different.

Enable the optional **revert-primary-port** option to direct the IPS port to switch back to the primary port when the primary port is up again. If this option is disabled (default) and the primary port is up again after a switchover, the old sessions will remain with the secondary port and do not switch back to the primary port. However, any new session will use the primary port. This is the only situation when both the primary and secondary ports are used at the same time.

LUN Trespass for Storage Port Failover


In addition to the high availability of statically imported iSCSI targets, the trespass feature is available to enable the move of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the trespass feature be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails,

Send documentation comments to dcnm-san-docfeedback@cisco.com

the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see Figure 4-13).

Figure 4-13 Virtual Target with an Active Primary Port

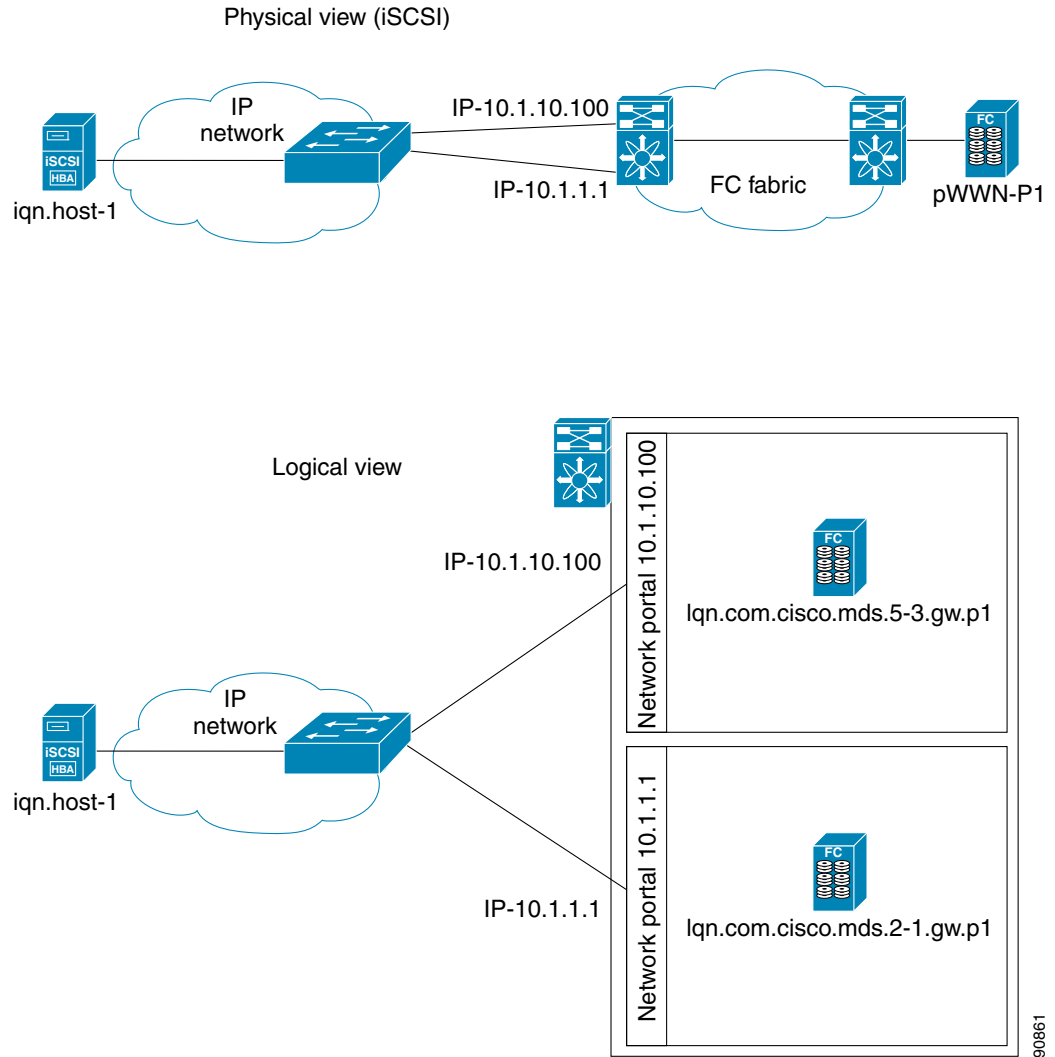


Multiple IPS Ports Connected to the Same IP Network

Figure 4-14 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-14 Multiple Gigabit Ethernet Interfaces in the Same IP Network



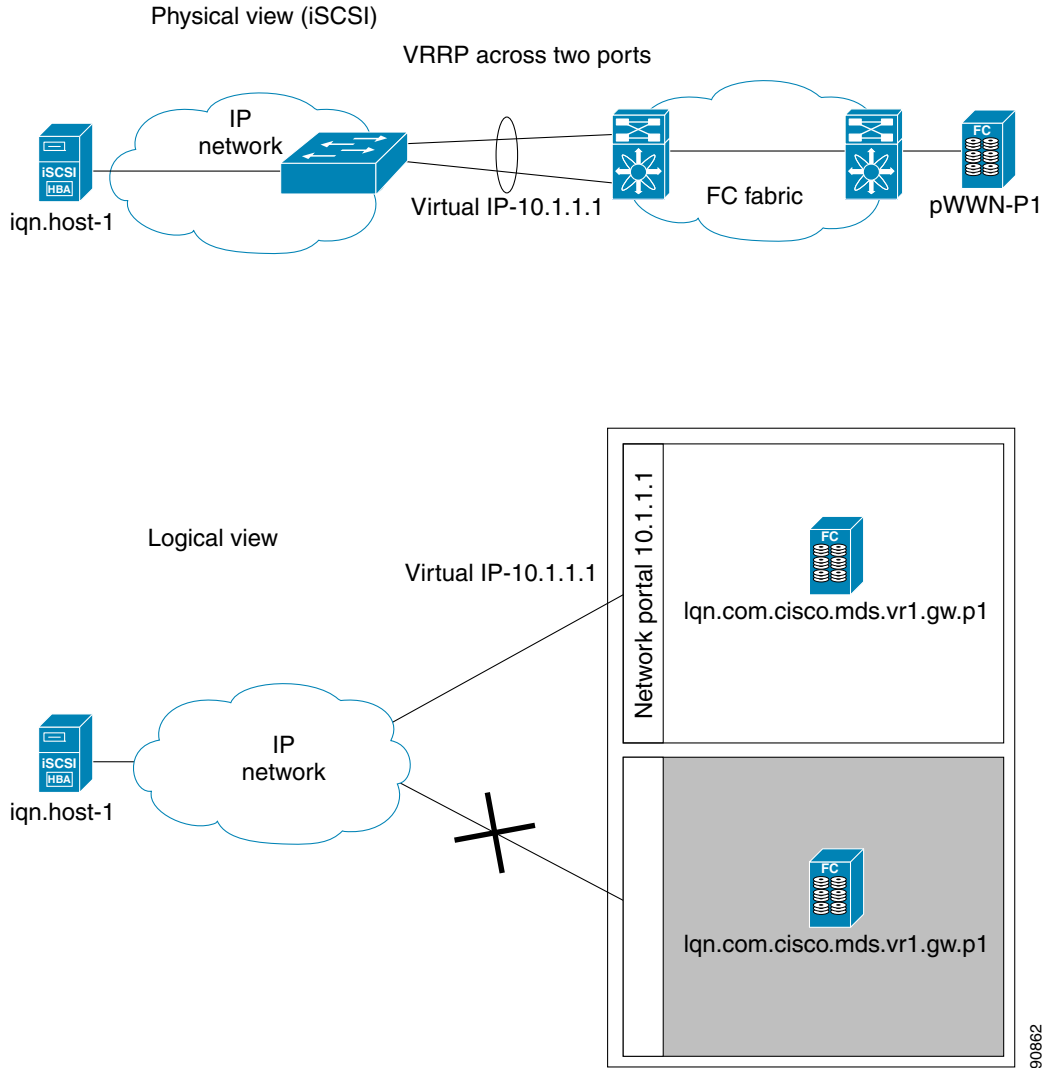
In Figure 4-14, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 4-15 provides an example of a VRRP-based high availability iSCSI configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-15 VRRP-Based iSCSI High Availability



In Figure 4-15, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Ethernet PortChannel-Based High Availability




Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that iSCSI link.

Figure 4-16 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-16 *Ethernet PortChannel-Based iSCSI High Availability*



In Figure 4-16, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.



Note

If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, the load balancing policy on the Ethernet switch must be based on source/destination IP address only, not port numbers, for load balancing with VRRP to operate correctly.

iSNS

Internet Storage Name Service (iSNS) allows your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. To facilitate these functions, the iSNS server and client function as follows:

- The iSNS client registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server.
- The iSNS server provides the following services for the iSNS client:
 - Device registration
 - State change notification
 - Remote domain discovery services

All iSCSI devices (both initiator and target) acting as iSNS clients, can register with an iSNS server. iSCSI initiators can then query the iSNS server for a list of targets. The iSNS server will respond with a list of targets that the querying client can access based on configured access control parameters.

A Cisco MDS 9000 Family switch can act as an iSNS client and register all available iSCSI targets with an external iSNS server. All switches in the Cisco MDS 9000 Family with IPS modules or MPS-14/2 modules installed support iSNS server functionality. This allows external iSNS clients, such as an iSCSI initiator, to register with the switch and discover all available iSCSI targets in the SAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About iSNS Client Functionality

Internet Storage Name Service (iSNS) allows your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. The iSNS client registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server. All iSCSI devices (both initiator and target) acting as iSNS clients can register with an iSNS server. When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server.

The iSNS client functionality on each IPS interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with an iSNS server.

Once a profile is tagged to an interface, the switch opens a TCP connection to the iSNS server IP address (using the well-known iSNS port number 3205) in the profile and registers network entity and portal objects; a unique entity is associated with each IPS interface. The switch then searches the Fibre Channel name server (FCNS) database and switch configuration to find storage nodes to register with the iSNS server.

Statically mapped virtual targets are registered if the associated Fibre Channel pWWN is present in the FCNS database and no access control configuration prevents it. A dynamically mapped target is registered if dynamic target importing is enabled. See the [“Presenting Fibre Channel Targets as iSCSI Targets” section on page 4-4](#) for more details on how iSCSI imports Fibre Channel targets.

A storage node is deregistered from the iSNS server when it becomes unavailable when a configuration changes (such as access control change or dynamic import disabling) or the Fibre Channel storage port goes offline. It is registered again when the node comes back online.

When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server. The iSNS client uses a registration interval value of 15 minutes. If the client fails to refresh the registration during this interval, the server will deregister the entries.

Untagging a profile also causes the network entity and portal to be deregistered from that interface.



Note

The iSNS client is not supported on a VRRP interface.

About iSNS Server Functionality

When enabled, the iSNS server on the Cisco 9000 Family MDS switch tracks all registered iSCSI devices. As a result, iSNS clients can locate other iSNS clients by querying the iSNS server. The iSNS server also provides the following functionalities:

- Allows iSNS clients to register, deregister, and query other iSNS clients registered with the iSNS server.
- Provides centralized management for enforcing access control to provide or deny access to targets from specific initiators.
- Provides a notification mechanism for registered iSNS clients to receive change notifications on the status change of other iSNS clients.
- Provides a single access control configuration for both Fibre Channel and iSCSI devices.
- Discovers iSCSI targets that do not have direct IP connectivity to the iSCSI initiators.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

iSNS Client Registration and Deregistration

You can use the **show isns database** command to display all registered iSNS clients and their associated configuration.

An iSNS client cannot query the iSNS server until it has registered. iSNS client deregistration can occur either explicitly or when the iSNS server detects that it can no longer reach the client (through ESI monitoring).

iSNS client registration and deregistration result in status change notifications (SCNs) being generated to all interested iSNS clients.

Target Discovery

iSCSI initiators discover targets by issuing queries to the iSNS server. The server supports *DevGetNext* requests to search the list of targets and *DevAttrQuery* to determine target and portal details, such as the IP address or port number to which to connect.

On receiving a query request from the iSCSI client, the iSNS server queries the Fibre Channel Name Server (FCNS) to obtain a list of Fibre Channel targets that are accessible by the querying initiator. The result of this query depends on zoning configuration currently active and current configuration(s) of the initiator. The iSNS server will subsequently use the iSCSI target configuration(s) (virtual target and dynamic import configuration) to translate the Fibre Channel target to an equivalent iSCSI target. At this stage it also applies any access control configured for the virtual target. A response message with the target details is then sent back to the query initiator.

The iSNS server sends a consolidated response containing all possible targets and portals to the querying initiator. For example, if a Fibre Channel target is exported as different iSCSI targets on different IPS interfaces, the iSNS server will respond with a list of all possible iSCSI targets and portals.

In order to keep the list of targets updated, the iSNS server sends state change notifications (SCN) to the client whenever an iSCSI target becomes reachable or unreachable. The client is then expected to rediscover its list of accessible targets by initiating another iSNS query. Reachability of iSCSI targets changes when any one of the following occurs:

- Target goes up or down.
- Dynamic import of FC target configuration changes.
- Zone set changes.
- Default zone access control changes.
- IPS interface state changes.
- Initiator configuration change makes the target accessible or inaccessible.

About Cloud Discovery

When an iSNS server receives a query request, it responds with a list of available targets and the portals through which the initiator can reach the target. The IP network configuration outside the MDS switch may result in only a subset of Gigabit Ethernet interfaces being reachable from the initiator. To ensure that the set of portals returned to the initiator is reachable, the iSNS server needs to know the set of Gigabit Ethernet interfaces that are reachable from a given initiator.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

iSNS Cloud Discovery is not supported on the Cisco Fabric Switch for IBM BladeCenter and Cisco Fabric Switch for HP c-Class BladeSystem.

The iSNS cloud discovery feature provides information to the iSNS server on the various interfaces reachable from an initiator by partitioning the interfaces on a switch into disjointed IP clouds. This discovery is achieved by sending messages to all other known IPS ports that are currently up and, depending on the response (or the lack of it), determines if the remote IPS port is in the same IP network or in a different IP network.

Cloud discovery is initiated when the following events occur:

- Manual requests from the CLI initiate cloud discovery from the CLI. This action causes the destruction of existing memberships and makes new ones.
- Auto-discovery of the interface results in an interface being assigned to its correct cloud. All other cloud members are not affected. The membership of each cloud is built incrementally and is initiated by the following events:
 - A Gigabit Ethernet interface comes up. This can be a local or remote Gigabit Ethernet interface.
 - The IP address of a Gigabit Ethernet interface changes.
 - The VRRP configuration on a port changes.

The iSNS server distributes cloud and membership information across all the switches using CFS. Therefore, the cloud membership view is the same on all the switches in the fabric.

**Note**

For CFS distribution to operate correctly for iSNS cloud discovery, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or NX-OS 4.1(1b) and later.

Licensing Requirements for iSCSI

The following table shows the licensing requirements for this feature:

License	License Description
Enterprise package (ENTERPRISE_PKG)	It comprises the IP security (IPsec) protocol for iSCSI and FCIP using the MPS-14/2 module or Cisco MDS 9216i Switch.

Guidelines and Limitations

iSLB configuration has the following limits:

- The maximum number of iSLB and iSCSI initiators supported in a fabric is 2000.
- The maximum number of iSLB and iSCSI sessions supported by an IPS port in either transparent or proxy initiator mode is 500.
- The maximum number of iSLB initiators supported in a fabric is 2000.
- The maximum number of iSLB initiators and iSCSI sessions supported by a switch is 5000.
- The maximum number of iSLB sessions per IPS port in either transparent or proxy initiator mode is 500.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- The maximum number of iSLB and iSCSI targets supported in a fabric is 6000.
- The maximum number of switches in a fabric that can have iSLB with CFS distribution enabled is four.
- No more than 200 new iSLB initiators can be added to the pending configuration. Before adding more initiators, you must commit the configuration.
- You cannot disable iSCSI if you have more than 200 iSLB initiators in the running configuration. Reduce the number of iSLB initiators to fewer than 200 before disabling iSCSI.
- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic is disrupted when any zoneset is activated.
- If IVR and iSLB features are enabled in the same fabric, you should have at least one switch in the fabric where both these features are enabled. Any zoning-related configuration and activation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, there may be traffic disruption in the fabric.

Default Settings

Table 4-2 lists the default settings for iSCSI parameters.

Table 4-2 **Default iSCSI Parameters**

Parameters	Default
Number of TCP connections	One per iSCSI session
minimum-retransmit-time	300 msec
keepalive-timeout	60 seconds
max-retransmissions	4 retransmissions
PMTU discovery	Enabled
pmtu-enable reset-timeout	3600 sec
SACK	Enabled
max-bandwidth	1 Gbps
min-available-bandwidth	70 Mbps
round-trip-time	1 msec
Buffer size	4096 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
Jitter	500 microseconds
TCP connection mode	Active mode is enabled
Fibre Channel targets to iSCSI	Not imported
Advertising iSCSI target	Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 4-2 *Default iSCSI Parameters (continued)*

Parameters	Default
iSCSI hosts mapping to virtual Fibre Channel hosts	Dynamic mapping
Dynamic iSCSI initiators	Members of the VSAN 1
Identifying initiators	iSCSI node names
Advertising static virtual targets	No initiators are allowed to access a virtual target (unless explicitly configured)
iSCSI login authentication	CHAP or none authentication mechanism
revert-primary-port	Disabled
Header and data digest	Enabled automatically when iSCSI initiators send requests. This feature cannot be configured and is not available in store-and-forward mode.
iSNS registration interval	60 sec (not configurable)
iSNS registration interval retries	3
Fabric distribution	Disabled

Table 4-3 lists the default settings for iSLB parameters.

Table 4-3 *Default iSLB Parameters*

Parameters	Default
Fabric distribution	Disabled
Load balancing metric	1000

Configuring iSCSI

This section describes how to configure iSCSI on the Cisco MDS 9000 Family switches.

This section includes the following sections:

- [Enabling iSCSI, page 4-35](#)
- [Creating iSCSI Interfaces, page 4-36](#)
- [Using the iSCSI Wizard, page 4-37](#)
- [Enabling Dynamic Mapping, page 4-37](#)
- [Creating Static Mapping, page 4-38](#)
- [Advertising Static iSCSI Targets, page 4-39](#)
- [Specifying the Initiator Identification, page 4-39](#)
- [Configuring the iSCSI Initiator Idle Timeout, page 4-40](#)
- [Configuring Dynamic Mapping, page 4-41](#)
- [Configuring Static Mapping, page 4-41](#)
- [Making the Dynamic iSCSI Initiator WWN Mapping Static, page 4-43](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Checking for WWN Conflicts, page 4-43](#)
- [Configuring the Proxy Initiator, page 4-44](#)
- [Configuring VSAN Membership for iSCSI Hosts, page 4-45](#)
- [Configuring Default Port VSAN for iSCSI Interfaces, page 4-46](#)
- [Adding iSCSI Initiator to the Zone Database, page 4-47](#)
- [Configuring Access Control in iSCSI, page 4-48](#)
- [Configuring AAA Authentication for an iSCSI User, page 4-50](#)
- [Configuring Authentication Mechanism, page 4-50](#)
- [Changing the iSCSI Routing Mode, page 4-55](#)
- [Configuring iSLB Using Device Manager, page 4-56](#)
- [Configuring iSLB Initiator Names or IP Addresses, page 4-57](#)
- [Assigning VSAN Membership for iSLB Initiators, page 4-59](#)
- [Configuring and Activating Zones for iSLB Initiators and Initiator Targets, page 4-63](#)
- [Restricting iSLB Initiator Authentication, page 4-64](#)

Enabling iSCSI

To use the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric. Alternatively, you can enable or disable the iSCSI feature directly on the required modules using Cisco DCNM for SAN or Device Manager. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.



Caution

When you disable this feature, all related configurations are automatically discarded.

Detailed Steps

To enable iSCSI on any participating switch, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters the configuration commands, one per line. End with CNTL/Z.
Step 2	switch(config)# feature iscsi	Enables iSCSI on that switch.
	switch(config)# iscsi enable module <x>	Enables iSCSI modules on the switch. Note New command added so that SME and iSCSI are available on the same switch.
	switch(config)# no iscsi enable module <x>	Disables the iSCSI module on the switch.
	switch(config)# no feature iscsi	Disables (default) iSCSI on that switch.

To enable iSCSI on any switch, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Choose **FC Services > iSCSI** in the Physical Attributes pane.

You see the iSCSI tables in the Information pane.

The **Control** tab is the default tab. You see the iSCSI enable status for all switches in the fabric that contain IPS ports.

Step 2 Choose **enable** from the Command column for each switch that you want to enable iSCSI on.

Step 3 Click the **Apply Changes** icon to save these changes.

To enable iSCSI on a module, follow these steps:

Step 1 Choose **FC Services > iSCSI** in the Physical Attributes pane.

You see the iSCSI tables in the Information pane.

Step 2 Click the **Module Control** tab.

You see the Module Control dialog box in the information pane.

Step 3 Check the **Mode Admin** check box to enable iSCSI for a specified port on the selected module.

Step 4 Click the **Apply Changes** icon to save these changes.

To enable iSCSI on a module using Device Manager, follow these steps:

Step 1 Choose **IP > iSCSI**

You see the iSCSI table.

Step 2 Check the **Mode Admin** check box to enable iSCSI for the specified port on the selected module.

Step 3 Click **Apply** to save these changes.

Creating iSCSI Interfaces

Each physical Gigabit Ethernet interface on an IPS module or MPS-14/2 module can be used to translate and route iSCSI requests to Fibre Channel targets and responses in the opposite direction. To enable this capability, the corresponding iSCSI interface must be in an enabled state.

Detailed Steps

To enable iSCSI interfaces, follow these steps:

Step 1 Enable the required Gigabit Ethernet interface.

```
switch# config terminal
switch(config)# interface gigabitethernet 2/1
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 2 Create the required iSCSI interface and enable the interface.

```
switch(config)# interface iscsi 2/1
switch(config-if)# no shutdown
```

Using the iSCSI Wizard

Detailed Steps

To use the iSCSI wizard in Cisco DCNM-SAN, follow these steps:

- Step 1** Click the **iSCSI Setup Wizard** icon.
You see the iSCSI Wizard Configure Initiator dialog box.
- Step 2** Select an existing iSCSI initiator or add the iSCSI node name or IP address for a new iSCSI initiator.
- Step 3** Select the switch for this iSCSI initiator if you are adding a new iSCSI initiator and click **Next**.
You see the iSCSI Wizard Select Targets dialog box.
- Step 4** Select the VSAN and targets to associate with this iSCSI initiator and click **Next**.



Note The iSCSI wizard turns on the Dynamic Import FC Targets feature.

You see the iSCSI Wizard Select Zone dialog box.

- Step 5** Set the zone name for this new iSCSI zone and check the **ReadOnly** check box if needed.
- Step 6** Click **Finish** to create this iSCSI initiator.
If created, the target VSAN is added to the iSCSI host VSAN list.



Note iSCSI wizard automatically turns on the Dynamic FC target import.

Enabling Dynamic Mapping

Detailed Steps

To enable dynamic mapping of Fibre Channel targets into iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi import target fc	IPS modules and MPS-14/2 modules dynamically import all Fibre Channel targets in the Fibre Channel SAN into the IP network.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To enable dynamic mapping of Fibre Channel targets into iSCSI using Device Manager, follow these steps:

-
- Step 1** Choose **IP > iSCSI**.
You see the iSCSI configuration.
 - Step 2** Click the **Target** tab to display a list of existing iSCSI targets.
 - Step 3** Check the **Dynamically Import FC Targets** check box.
 - Step 4** Click **Apply** to save this change.
-

Creating Static Mapping

Detailed Steps

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

-
- Step 1** Click **IP > iSCSI**.
You see the iSCSI configuration.
 - Step 2** Click the **Targets** tab to display a list of existing iSCSI targets .
 - Step 3** Click **Create** to create an iSCSI target.
You see the Create iSCSI Targets dialog box.
 - Step 4** Set the iSCSI target node name in the iSCSI Name field, in IQN format.
 - Step 5** Set the Port WWN field for the Fibre Channel target port you are mapping.
 - Step 6** Click the **Select from List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or click the **All** radio button to let the iSCSI target access all iSCSI initiators. Also see the “iSCSI Access Control” section on page 4-11.
 - Step 7** Click the **Select from List** radio button and check each interface you want to advertise the iSCSI targets on or click the **All** radio button to advertise all interfaces.
 - Step 8** Click **Apply** to save this change.
-



Tip

An iSCSI target cannot contain more than one Fibre Channel target port. If you have already mapped the whole Fibre Channel target port, you cannot use the LUN mapping option.



Note

See the “iSCSI-Based Access Control” section on page 4-12 for more information on controlling access to statically mapped targets.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Advertising Static iSCSI Targets

You can limit the Gigabit Ethernet interfaces through which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

Detailed Steps

To configure a specific interface that should advertise the iSCSI virtual target using Device Manager, follow these steps:

-
- Step 1** Select **IP > iSCSI**.
You see the iSCSI configuration.
- Step 2** Click the **Targets** tab to display a list of existing iSCSI targets.
- Step 3** Right-click the iSCSI target that you want to modify and click **Edit Advertised**.
You see the Advertised Interfaces dialog box.
- Step 4** (Optional) Right-click an interface that you want to delete and click **Delete**.
- Step 5** (Optional) Click **Create** to advertise on more interfaces.
You see the Create Advertised Interfaces dialog box.
-

To configure a specific interface that should advertise the iSCSI virtual target, 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 or MPS-14/2 modules. Note To advertise the virtual target on multiple interfaces, issue the command for each interface.
	<code>switch(config-iscsi-tgt)# no advertise interface GigabitEthernet 2/5</code>	Removes this interface from the list of interfaces from which this target is advertised.

Specifying the Initiator Identification

You can configure the iSCSI initiator identification mode on each IPS port and all the iSCSI hosts terminating on the IPS port will be identified according to that configuration. The default mode is to identify the initiator by name.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To specify the initiator identification mode, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters the configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	switch(config-if)# switchport initiator id ip-address	Identifies the iSCSI initiator based on the IP address.
	switch(config-if)# switchport initiator id name	Identifies the iSCSI initiator based on the initiator node name. This is the default behavior.

To specify the initiator identification mode, follow these steps:

-
- Step 1** Choose **Interfaces > FC Logical** from the Physical Attributes pane.
You see the interfaces configuration in the Information pane.
- Step 2** Click the **iSCSI** tab.
You see the iSCSI interfaces configuration.
- Step 3** Right-click the Initiator ID Mode field for the iSCSI interface that you want to modify and select **name** or **ipaddress** from the drop-down menu.
- Step 4** Click **Apply Changes** to save this change.
-

Configuring the iSCSI Initiator Idle Timeout

iSCSI initiator idle timeout specifies the time for which the virtual Fibre Channel N port is kept idle after the initiator logs out from its last iSCSI session. The default value for this timer is 300 seconds. This is useful to avoid N ports logging in to and logging off of the Fibre Channel SAN as transient failure occurs in the IP network. This helps reduce unnecessary RSCNs being generated in the Fibre Channel SAN.

Detailed Steps

To configure the initiator idle timeout, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters the configuration mode.
Step 2	switch(config)# iscsi initiator idle-timeout 10	Configures the iSCSI initiators to have an idle timeout value of 10 seconds.

To configure the initiator idle timeout, follow these steps:

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Click the **Globals** tab.
You see the iSCSI global configuration.
- Step 3** Right-click on the InitiatorIdle Timeout field that you want to modify and enter the new timeout value.
- Step 4** Click the **Apply Changes** icon to save these changes.

Configuring Dynamic Mapping

Detailed Steps

To configure dynamic mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi dynamic initiator islb	Specifies iSLB dynamic initiator mode.
	switch(config)# iscsi dynamic initiator deny	Disallows dynamic initiators from logging on to the MDS switch.
	switch(config)# no iscsi dynamic initiator islb	Reverts to iSCSI mode (default).

Configuring Static Mapping

Detailed Steps

To configure static mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)#	Configures an iSCSI initiator using the iSCSI name of the initiator node. The maximum name length is restricted to 223 alphanumeric characters. The minimum length is 16.
	switch(config)# no iscsi initiator name iqn.1987-02.com.cisco.initiator	Deletes the configured iSCSI initiator.

To configure static mapping for an iSCSI initiator using Device Manager, follow these steps:

- Step 1** Select **IP > iSCSI**.
You see the iSCSI configuration. The Initiators tab is the default.
- Step 2** Click **Create** to create an iSCSI initiator.
You see the Create iSCSI Initiators dialog box.
- Step 3** Set the iSCSI node name or IP address and VSAN membership.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 4** In the Node WWN section, check the **Persistent** check box.
- Step 5** Check the **System Assigned** check box if you want the switch to assign the nWWN or leave this unchecked and set the Static WWN field.
- Step 6** In the Port WWN section, check the **Persistent** check box if you want to statically map pWWNs to the iSCSI initiator.
- Step 7** If persistent, check the **System Assigned** check box and set the number of pWWNs to reserve for this iSCSI initiator if you want the switch to assign pWWNs. Alternately, you can leave this unchecked and set one or more pWWNs for this iSCSI initiator.
- Step 8** (Optional) Set the AuthUser field if authentication is enabled. Also see the “iSCSI Session Authentication” section on page 4-13.
- Step 9** Click **Create** to create this iSCSI initiator.

To configure static mapping (using the **ip-address** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator ip-address 10.50.0.0 switch(config-iscsi-init)#	Configures an iSCSI initiator using the IPv4 address of the initiator node.
	switch(config)# iscsi initiator ip-address 2001:0DB8:800:200C::417A switch(config-iscsi-init)#	Configures an iSCSI initiator using the IPv6 unicast address of the initiator node.
	switch(config)# no iscsi initiator ip-address 2001:0DB8:800:200C::417A	Deletes the configured iSCSI initiator.

To assign the WWN for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch(config-iscsi-init)# static nwwn system-assign	Uses the switch’s WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent.
	switch(config-iscsi-init)# static nwwn 20:00:00:05:30:00:59:11	Assigns the user provided WWN as the nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node.
Step 2	switch(config-iscsi-init)# static pwwn system-assign 2	Uses the switch’s WWN pool to allocate two pWWNs for this iSCSI initiator and keeps them persistent. The range is from 1 to 64.
	switch(config-iscsi-init)# static pwwn 21:00:00:20:37:73:3b:20	Assigns the user provided WWN as the pWWN for the iSCSI initiator.



Note

If the system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is saved to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a write erase, you must manually delete the WWN configuration from the ASCII file. Failing to do so can cause duplicate WWN assignments if the ASCII configuration file is reapplied on the switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Making the Dynamic iSCSI Initiator WWN Mapping Static

After a dynamic iSCSI initiator has already logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping so this initiator uses the same mapping the next time it logs in.

You can convert a dynamic iSCSI initiator to static iSCSI initiator and make its WWNs persistent (see the “Dynamic Mapping” section on page 4-9).



Note

You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator.



Note

Making the dynamic pWWNs static after the initiator is created is supported only through the CLI, not through Device Manager or Cisco DCNM- SAN. In Cisco DCNM-SAN or Device Manager, you must delete and then recreate this initiator to have the pWWNs static.

Detailed Steps

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi save-initiator name iqn.1987-02.com.cisco.initiator	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose name is specified.
	switch(config)# iscsi save-initiator ip-address 10.10.100.11	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv4 address is specified.
	switch(config)# iscsi save-initiator ip-address 2001:0DB8:800:200C::417A	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv6 unicast address is specified.
	switch(config)# iscsi save-initiator	Saves the nWWN and pWWNs that have automatically been assigned to all the initiators.
Step 3	switch(config)# exit switch#	Returns to EXEC mode.
Step 4	switch# copy running-config startup-config	Saves the nWWN/pWWN mapping configuration across system reboots.

Checking for WWN Conflicts

WWNs assigned to static iSCSI initiators by the system can be inadvertently returned to the system when an upgrade fails or you downgrade the system software (manually booting up an older Cisco MDS SAN-OS release without using the **install all** command). In these instances, the system can later assign those WWNs to other iSCSI initiators (dynamic or static) and cause conflicts.

You can address this problem by checking for and removing any configured WWNs that belong to the system whenever such scenarios occur.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To check for and remove WWN conflicts, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi duplicate-wwn-check List of Potential WWN Conflicts: ----- Node : iqn.test-local-nwwn:1-local-pwwn:1 nWWN : 22:03:00:0d:ec:02:cb:02 pWWN : 22:04:00:0d:ec:02:cb:02	Checks for WWN conflicts.
Step 3	switch(config)# iscsi initiator name iqn.test-local-nwwn:1-local-pwwn:1	Enters iSCSI initiator configuration mode for the initiator named iqn.test-local-nwwn:1-local-pwwn:1.
Step 4	switch(config-iscsi-init)# no static nWWN 22:03:00:0d:ec:02:cb:02	Removes a conflicting nWWN.
Step 5	switch(config-iscsi-init)# no static pWWN 22:04:00:0d:ec:02:cb:02	Removes a conflicting pWWN.

To permanently keep the automatically assigned nWWN mapping, follow these steps:

-
- Step 1 Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2 Click the **Initiators** tab.
You see the iSCSI initiators configured.
- Step 3 Check the **Persistent Node WWN** check box for the iSCSI initiators that you want to make static.
- Step 4 Click the **Apply Changes** icon to save these changes.
-

Configuring the Proxy Initiator

Detailed Steps

To configure the proxy initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that initiators will connect to.
Step 3	switch(config-if)# switchport proxy-initiator	Configures the proxy initiator mode with system-assignment nWWN and pWWN.
	switch(config-if)# no switchport proxy-initiator	Disables the proxy initiator mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	<code>switch(config-if)# switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22</code>	(Optional) Configures the proxy initiator mode using the specified WWNs.
	<code>switch(config-if)# no switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22</code>	Disables the proxy initiator mode.

To configure the proxy initiator, follow these steps:

-
- Step 1** Expand **Switches**, expand **FC Interfaces**, and then select **Logical** in the Physical Attributes pane. You see the Interface tables in the Information pane.
 - Step 2** In Device Manager, select **Interface > Ethernet and iSCSI**. You see the Ethernet Interfaces and iSCSI dialog box.
 - Step 3** Click the **iSCSI** tab in either FM or DM. You see the iSCSI interface configuration table.
 - Step 4** Check the **Proxy Mode Enable** check box.
 - Step 5** Click the **Apply Changes** icon in Cisco DCNM-SAN or click **Apply** in Device Manager to save these changes.
-



Note

When an interface is in proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the iSCSI interface's proxy N port attributes—the WWN pairs or the FC ID. You cannot configure zoning using iSCSI attributes such as IP address or IQN of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the [“iSCSI Access Control” section on page 4-11](#)).

Configuring VSAN Membership for iSCSI Hosts

Individual iSCSI hosts can be configured to be in a specific VSAN. The specified VSAN overrides the iSCSI interface VSAN membership.

Detailed Steps

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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
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.

To assign VSAN membership for iSCSI hosts, follow these steps:

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Click the **Initiators** tab.
You see the iSCSI initiators configured.
- Step 3** Fill in the VSAN Membership field to assign a VSAN to the iSCSI hosts.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

**Note**

When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

Configuring Default Port VSAN for iSCSI Interfaces

VSAN membership can be configured for an iSCSI interface, called the *port VSAN*. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. In other words, the port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1.

**Caution**

Changing the VSAN membership of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [“Changing iSCSI Interface Parameters and the Impact on Load Balancing”](#) section on page 4-21.

Detailed Steps

To change the default port VSAN for an iSCSI interface, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# iscsi interface</code> <code>vsan-membership</code>	Enables you to configure VSAN membership for iSCSI interfaces.
Step 3	<code>switch(config)# vsan database</code> <code>switch(config-vsan-db)#</code>	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	<code>switch(config-vsan-db)# vsan 2 interface iscsi 2/1</code>	Assigns the membership of the iscsi 2/1 interface to the specified VSAN (VSAN 2).
	<code>switch(config-vsan-db)# no vsan 2 interface iscsi 2/1</code>	Reverts to using the default VSAN as the port VSAN of the iSCSI interface.

To change the default port VSAN for an iSCSI interface using Device Manager, follow these steps:

-
- Step 1** Choose [Interface > Ethernet and iSCSI](#).
You see the Ethernet Interfaces and iSCSI dialog box.
- Step 2** Click the [iSCSI](#) tab.
You see the iSCSI interface configuration table.
- Step 3** Double-click the PortVSAN column and modify the default port VSAN.
- Step 4** Click [Apply](#) to save these changes.
-

Adding iSCSI Initiator to the Zone Database

Detailed Steps

To add an iSCSI initiator to the zone database, follow these steps:

	Command	Purpose
Step 1	<code>switch# config terminal switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# zone name iSCSIzone vsan 1 switch(config-zone)</code>	Creates a zone name for the iSCSI devices in the IPS module or MPS-14/2 module to be included.
Step 3	<code>switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1</code>	Assigns an iSCSI node name-based membership into a zone.
	<code>switch(config-zone)# no member symbolic-nodename iqn.1987-02.com.cisco.init1</code>	Deletes the specified device from a zone.
	<code>switch(config-zone)# member ip-address 10.50.1.1</code>	Assigns an iSCSI IPv4 address-based membership into a zone.
	<code>switch(config-zone)# no member ip-address 10.50.1.1</code>	Deletes the specified device from a zone.
	<code>switch(config-zone)# member ipv6-address 2001:0DB8:800:200C::417A</code>	Assigns an iSCSI IPv6 address-based membership into a zone.
	<code>switch(config-zone)# no member ipv6-address 2001:0DB8:800:200C::417A</code>	Deletes the specified device from a zone.
	<code>switch(config-zone)# member pwwn 20:00:00:05:30:00:59:11</code>	Assigns an iSCSI port WWN-based membership into a zone.
	<code>switch(config-zone)# no member pwwn 20:00:00:05:30:00:59:11</code>	Deletes the device identified by the port WWN from a zone.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To add an iSCSI initiator to the zone database, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Edit Local Zone Database dialog box.
- Step 2** Select the VSAN you want to add the iSCSI host initiator to and click **OK**.
You see the available zones and zone sets for that VSAN.
- Step 3** From the list of available devices with iSCSI host initiators, drag the initiators to add into the zone.
- Step 4** Click **Distribute** to distribute the change.
-

Configuring Access Control in iSCSI

Detailed Steps

To configure access control in iSCSI follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator switch(config-iscsi-tgt)#	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	<code>switch(config-iscsi-tgt)# initiator iqn.1987-02.com.cisco.initiator1 permit</code>	Allows the specified iSCSI initiator node to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	<code>switch(config-iscsi-tgt)# no initiator iqn.1987-02.com.cisco.initiator1 permit</code>	Prevents the specified initiator node from accessing virtual targets.
	<code>switch(config-iscsi-tgt)# initiator ip address 10.50.1.1 permit</code>	Allows the specified IPv4 address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	<code>switch(config-iscsi-tgt)# no initiator ip address 10.50.1.1 permit</code>	Prevents the specified IPv4 address from accessing virtual targets.
	<code>switch(config-iscsi-tgt)# initiator ip address 10.50.1.0 255.255.255.0 permit</code>	Allows all initiators in this IPv4 subnetwork (10.50.1/24) to access this virtual target.
	<code>switch(config-iscsi-tgt)# no initiator ip address 10.50.1.0 255.255.255.0 permit</code>	Prevents all initiators in this IPv4 subnetwork from accessing virtual targets.
	<code>switch(config-iscsi-tgt)# initiator ip address 2001:0DB8:800:200C::417A permit</code>	Allows the specified IPv6 unicast address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
	<code>switch(config-iscsi-tgt)# no initiator ip address 2001:0DB8:800:200C::417A permit</code>	Prevents the specified IPv6 address from accessing virtual targets.
	<code>switch(config-iscsi-tgt)# initiator ip address 2001:0DB8:800:200C::/64 permit</code>	Allows all initiators in this IPv6 subnetwork (2001:0DB8:800:200C::/64) to access this virtual target.
	<code>switch(config-iscsi-tgt)# no initiator ip address 2001:0DB8:800:200C::/64 permit</code>	Prevents all initiators in this IPv6 subnetwork from accessing virtual targets.
	<code>switch(config-iscsi-tgt)# all-initiator-permit</code>	Allows all initiator nodes to access this virtual target.
	<code>switch(config-iscsi-tgt)# no all-initiator-permit</code>	Prevents any initiator from accessing virtual targets (default).

To configure access control in iSCSI using Device Manager, follow these steps:

-
- Step 1** Select **IP > iSCSI**.
You see the iSCSI configuration.
 - Step 2** Click the **Targets** tab.
You see the iSCSI virtual targets.
 - Step 3** Uncheck the **Initiators Access All** check box if checked.
 - Step 4** Click **Edit Access**.
You see the Initiators Access dialog box.
 - Step 5** Click **Create** to add more initiators to the Initiator Access list.
You see the Create Initiators Access dialog box.
 - Step 6** Add the name or IP address for the initiator that you want to permit for this virtual target.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 7 Click **Create** to add this initiator to the Initiator Access List.

Configuring AAA Authentication for an iSCSI User

Detailed Steps

To configure AAA authentication for an iSCSI user, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# aaa authentication iscsi default group RadServerGrp	Uses RADIUS servers that are added in the group called RadServerGrp for the iSCSI CHAP authentication.
	switch(config)# aaa authentication iscsi default group TacServerGrp	Uses TACACS+ servers that are added in the group called TacServerGrp for the iSCSI CHAP authentication.
	switch(config)# aaa authentication iscsi default local	Uses the local password database for iSCSI CHAP authentication.

To configure AAA authentication for an iSCSI user, follow these steps:

Step 1 Choose **Switches > Security > AAA** in the Physical Attributes pane.

You see the AAA configuration in the Information pane.

Step 2 Click the **Applications** tab.

You see the AAA configuration per application.

Step 3 Right-click the ServerGroup Id List field for the iSCSI application and enter the server group that you want iSCSI to use.



Note You should use an existing server group or create a new server group before configuring it for iSCSI session authentication.

Step 4 Click the **Apply Changes** icon to save these changes.

Configuring Authentication Mechanism

You can configure iSCSI CHAP or None authentication at both the global level and at each interface level.

The authentication for a Gigabit Ethernet interface or subinterface overrides the authentication method configured at the global level.

If CHAP authentication is used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used at all, issue the **iscsi authentication none** command.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To configure the authentication mechanism for iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi authentication chap	Configures CHAP as the default authentication mechanism globally for the Cisco MDS switch. CHAP authentication is required for all iSCSI sessions.

To configure AAA authentication for an iSCSI user, follow these steps:

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
 - Step 2** Click the **Globals** tab.
You see the iSCSI authentication configuration table.
 - Step 3** Select **chap** or **none** from the authMethod column.
 - Step 4** Click the **Apply Changes** icon in Cisco DCNM-SAN to save these changes.
-

To configure the authentication mechanism for iSCSI sessions to a particular interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface GigabitEthernet 2/1.100 switch(config-if)#	Selects the Gigabit Ethernet interface.
Step 3	switch(config-if)# iscsi authentication none	Specifies that no authentication is required for iSCSI sessions to the selected interface.

To configure the authentication mechanism for iSCSI sessions to a particular interface, follow these steps:

-
- Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
You see the Gigabit Ethernet configuration in the Information pane.
 - Step 2** Click the **iSNS** tab.
You see the iSCSI and iSNS configuration.
 - Step 3** Right-click on the **IscsiAuthMethod** field and select none or chap.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Local Authentication

See the *Security Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS Security Guide to create the local password database. To create users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory.

Detailed Steps

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 ffsffsfffs345353554535 iscsi	Configures a user name (iscsiuser) and password (ffsffsfffs345353554535) in the local database for iSCSI login authentication.

To configure iSCSI users for local authentication using Device Manager, follow these steps:

-
- Step 1** Choose **Security > iSCSI**.
You see the iSCSI Security dialog box.
- Step 2** Complete the iSCSI User, Password, and Password Confirmation fields.
- Step 3** Click **Create** to save this new user.
-

Restricting iSCSI Initiator Authentication

By default, the iSCSI initiator can use any user name in the RADIUS server or in the local database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSCSI initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password has been compromised.

Detailed Steps

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.init switch(config-iscsi-init)#	Enters the configuration submode for the initiator iqn.1987-02.com.cisco.init.
Step 3	switch(config-iscsi-init)# username user1	Restricts the initiator iqn.1987-02.com.cisco.init to only authenticate using user1 as its CHAP user name. Tip Be sure to define user1 as an iSCSI user in the local AAA database or the RADIUS server.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Right-click the AuthUser field and enter the user name to which you want to restrict the iSCSI initiator.
- Step 3** Click the **Apply Changes** icon to save these changes.
-

Configuring Mutual CHAP Authentication

The IPS module or MPS-14/2 module supports a mechanism by which the iSCSI initiator can authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication is available in addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator.

In addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator, the IPS module or MPS-14/2 module also supports a mechanism for the iSCSI initiator to authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSCSI initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

Detailed Steps

To configure a global iSCSI target user name and password to be used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi authentication username testuser password abc123	Configures the switch user account (testuser) along with a password (abc123) specified in clear text (default) for all initiators. The password is limited to 128 characters.
	switch(config)# iscsi authentication username user1 password 7 !@*asdsfsdfjh!@df	Configures the switch user account (user1) along with the encrypted password specified by 7 (!@*asdsfsdfjh!@df) for all initiators.
	switch(config)# iscsi authentication username user1 password 0 abcd12AAA	Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (indicated by 0—default) for all initiators. The password is limited to 128 characters.
	switch(config)# no iscsi authentication username testuser	Removes the global configuration for all initiators.

To configure a global iSCSI target user name and password to be used by the switch to authenticate itself to an initiator, follow these steps:

-
- Step 1** Choose **FC Interfaces > Logical > iSCSI** in the Physical Attributes pane.
You see the iSCSI tables in the Information pane.
- Step 2** Select the **Globals** tab.
You see the global iSCSI configuration.
- Step 3** Fill in the Target UserName and Target Password fields.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click the **Apply Changes** icon to save these changes.

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)#	Configures an iSCSI initiator using the iSCSI name of the initiator node.
Step 3	switch(config-iscsi-init)# mutual-chap username testuser password abcd12AAA	Configures the switch user account (testuser) along with a password (abcd12AAA) specified in clear text (default). The password is limited to 128 characters.
	switch(config-iscsi-init)# mutual-chap username user1 password 7 !@*asdfsdfjh!@df	Configures the switch user account (user1) along with the encrypted password specified by 7 (!@*asdfsdfjh!@df).
	switch(config-iscsi-init)# no mutual-chap username testuser	Removes the switch authentication configuration.

Use the **show running-config** and the **show iscsi global** commands to display the global configuration. Use the **show running-config** and the **show iscsi initiator configured** commands to display the initiator specific configuration. (See the “Displaying iSCSI User Information” section on page 4-95 for command output examples).

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator using Device Manager, follow these steps:

Step 1 Choose **IP > iSCSI**.

You see the iSCSI configuration.

Step 2 Complete the Target UserName and Target Password fields for the initiator that you want to configure.

Step 3 Click **Create** to add this initiator to the Initiator Access List.

Configuring an iSCSI RADIUS Server

Detailed Steps

To configure an iSCSI RADIUS server, follow these steps:

Step 1 Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.

Step 2 Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.

Step 3 Configure the iSCSI users and passwords on the RADIUS server.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Setting QoS Values

Detailed Steps

To set the QoS values, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# qos 3</code>	Configures the differentiated services code point (DSCP) value of 3 to be applied to all outgoing IP packets in this iSCSI interface. The valid range for the iSCSI DSCP value is from 0 to 63.
Step 2	<code>switch(config-if)# no qos 5</code>	Reverts the switch to its factory default (marks all packets with DSCP value 0).


To set the QoS values, follow these steps:

-
- Step 1** Expand **Switches**, expand **FC Interfaces**, and then select **Logical** in the Physical Attributes pane. You see the Interface tables in the Information pane.
 - Step 2** In Device Manager, choose **Interface > Ethernet and iSCSI**. You see the Ethernet Interfaces and iSCSI dialog box.
 - Step 3** Click the **iSCSI TCP** tab in either Cisco DCNM-SAN or Device Manager. You see the iSCSI TCP configuration table.
 - Step 4** Set the QoS field from 1 to 6.
 - Step 5** Click the **Apply Changes** icon in Cisco DCNM-SAN or click **Apply** in Device Manager to save these changes.
-

Changing the iSCSI Routing Mode

Detailed Steps

To set the iSCSI routing mode, follow this step:

	Command	Purpose
Step 1	<code>switch(config-if)# mode cut-thru</code>	Configures cut-thru mode on the iSCSI interface.
		 Caution Changing the iSCSI routing mode disrupts the iSCSI sessions on the interface.
	<code>switch(config-if)# no mode cut-thru</code>	Reverts store-and-forward mode (default).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring iSLB



Note

For iSLB, all switches in the fabric must be running Cisco MDS SAN-OS Release 2.1(1a) or later.

This section covers the following topics:

- [Configuring iSLB Using Device Manager, page 4-56](#)
- [Configuring iSLB Initiator Names or IP Addresses, page 4-57](#)
- [Making the Dynamic iSLB Initiator WWN Mapping Static, page 4-59](#)
- [Assigning VSAN Membership for iSLB Initiators, page 4-59](#)
- [Configuring and Activating Zones for iSLB Initiators and Initiator Targets, page 4-63](#)
- [Restricting iSLB Initiator Authentication, page 4-64](#)
- [Mutual CHAP Authentication, page 4-65](#)
- [Configuring Load Balancing Using VRRP, page 4-66](#)
- [Enabling VRRP for Load Balancing, page 4-66](#)

Configuring iSLB Using Device Manager

Prerequisites

Perform the following actions prior to configuring iSLB:

- [Enable iSCSI](#) (see the “Enabling iSCSI” section on page 4-35 for more information).
- [Configure the Gigabit Ethernet interfaces](#) (see the “Configuring Gigabit Ethernet Interface” section on page 7-5).
- [Configure the VRRP groups](#) (see the “Configuring Load Balancing Using VRRP” section on page 4-66).
- [Configure and activate a zone set](#) (see the *Fabric Configuration Guide, Cisco DCNM for SAN Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* for more information).
- [Enable CFS distribution for iSLB](#) (see the “Enabling iSLB Configuration Distribution” section on page 4-67).

Detailed Steps

To configure iSLB using Device Manager, follow these steps:

-
- Step 1** Choose **IP > iSCSI iSLB**.
You see the iSCSI iSLB dialog box.
 - Step 2** Click **Create** to create a new iSCSI iSLB initiator.
You see the Create iSCSI iSLB Initiators dialog box.
 - Step 3** Set the Name or IP Address field to the iSLB name or IP address.
 - Step 4** Set the VSAN Membership field to the VSAN that you want the iSLB initiator in.
Also see the “Assigning VSAN Membership for iSLB Initiators” section on page 4-59.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 5** Check the **Persistent** check box to convert a dynamic nWWN to static for the iSLB initiator. Also see the “Making the Dynamic iSCSI Initiator WWN Mapping Static” section on page 4-43.
- Step 6** (Optional) Check the **SystemAssigned** check box to have the switch assign the nWWN.
- Step 7** (Optional) Set the Static WWN field to manually assign the static nWWN. You must ensure uniqueness for this nWWN.
- Step 8** (Optional) Check the Port WWN Mapping **Persistent** check box to convert dynamic pWWNs to static for the iSLB initiator. See the “Making the Dynamic iSCSI Initiator WWN Mapping Static” section on page 4-43.
- Step 9** (Optional) Check the **SystemAssigned** check box and set the number of pWWNs you want to have the switch assign the PWWN.
- Step 10** (Optional) Set the Static WWN(s) field to manually assign the static pWWNs. You must ensure uniqueness for these pWWN.
- Step 11** (Optional) Set the AuthUser field to the username that you want to restrict the iSLB initiator to for iSLB authentication. Also see the “Restricting iSLB Initiator Authentication” section on page 4-64.
- Step 12** Fill in the Username and Password fields to configure iSLB initiator target CHAP authentication. Also see the “iSLB Session Authentication” section on page 4-19.
- Step 13** In the Initiator Specific Target section, set the pWWN to configure an iSLB initiator target.
- Step 14** (Optional) Set the Name field to a globally unique identifier (IQN).
- Step 15** (Optional) Check the **NoAutoZoneCreation** check box to disable auto-zoning.
- Step 16** (Optional) Check the **TresspassMode** check box. Also see the “LUN Trespass for Storage Port Failover” section on page 4-25.
- Step 17** (Optional) Check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up.
- Step 18** Set the PrimaryVsan to the VSAN for the iSLB initiator target.
- Step 19** Click **Create** to create this iSLB initiator.
- Step 20** If CFS is enabled, select **commit** from the CFS drop-down menu.

Configuring iSLB Initiator Names or IP Addresses

You must specify the iSLB initiator name or IP address before configuring it.



Note

Specifying the iSLB initiator name or IP address is the same as for an iSCSI initiator. See the “Static Mapping” section on page 4-9.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To enter iSLB initiator configuration submode using the **name** option for an iSLB initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator switch(config-islb-init)#	Configures an iSLB initiator using the iSCSI name of the initiator node (iqn.1987-02.com.cisco.initiator) and enters iSLB initiator configuration submode. The maximum name length is 223 alphanumeric characters. The minimum length is 16.
	switch(config)# no islb initiator name iqn.1987-02.com.cisco.initiator	Deletes the configured iSLB initiator.

To enter iSLB initiator configuration submode using the **ip-address** option for an iSLB initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator ip-address 10.1.1.3 switch(config-islb-init)#	Configures an iSLB initiator using the IPv4 address of the initiator node and enters iSLB initiator configuration submode.
	switch(config)# no islb initiator ip-address 10.1.1.3	Deletes the configured iSLB initiator.
	switch(config)# islb initiator ip-address 2001:0DB8:800:200C::417A switch(config-islb-init)#	Configures an iSLB initiator using the IPv6 unicast address of the initiator node and enters iSLB initiator configuration submode.
	switch(config)# no islb initiator ip-address 2001:0DB8:800:200C::417A	Deletes the configured iSLB initiator.

Examples

To verify the iSLB initiator configuration, use the **show islb initiator configured** command.

```
switch# show islb initiator configured
iSCSI Node name is 10.1.1.2
  Member of vsans: 10
  Node WWN is 23:02:00:0c:85:90:3e:82
  Load Balance Metric: 100
  Number of Initiator Targets: 1

  Initiator Target: test-target
    Port WWN 01:01:01:01:02:02:02:02
    Primary PWWN VSAN 1
    Zoning support is enabled
    Trespass support is disabled
    Revert to primary support is disabled
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Making the Dynamic iSLB Initiator WWN Mapping Static

After a dynamic iSLB initiator has logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping to allow this initiator to use the same mapping the next time it logs in (see the “Dynamic Mapping” section on page 4-5).

You can convert a dynamic iSLB initiator to a static iSLB initiator and make its WWNs persistent.

Restrictions

- You cannot convert a dynamic iSCSI initiator to a static iSLB initiator or a dynamic iSLB initiator to a static iSCSI initiator (see the “Dynamic Mapping” section on page 4-20).
- Making the dynamic mapping for iSLB initiators static is the same as for iSCSI. See the “Making the Dynamic iSCSI Initiator WWN Mapping Static” section on page 4-43.
- Only statically mapped iSLB initiator configuration is distributed throughout the fabric using CFS. Dynamically and statically configured iSCSI initiator configurations are not distributed.

See the “Configuring iSLB Using Device Manager” procedure on page 4-56.

Detailed Steps

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb save-initiator name iqn.1987-02.com.cisco.initiator	Saves the nWWNs and pWWNs that have automatically been assigned to the iSLB initiator whose name is specified.
	switch(config)# islb save-initiator 10.10.100.11	Saves the nWWNs and pWWNs that have automatically been assigned to the iSLB initiator whose IPv4 address is specified.
	switch(config)# iscsi save-initiator ip-address 2001:0DB8:800:200C::417A	Saves the nWWNs and pWWNs that have automatically been assigned to the iSCSI initiator whose IPv6 unicast address is specified.
	switch(config)# islb save-initiator	Saves the nWWNs and pWWNs that have automatically been assigned to all the iSLB initiators.
Step 3	switch(config)# exit switch#	Returns to EXEC mode.
Step 4	switch# copy running-config startup-config	Saves the nWWN/pWWN mapping configuration across system reboots.

Assigning VSAN Membership for iSLB Initiators

Individual iSLB hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel). The specified VSAN overrides the iSCSI interface VSAN membership.

For more information, see the *Fabric Configuration Guide, Cisco DCNM for SAN*.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

Specifying the iSLB initiator VSAN is the same as for an iSCSI initiator. See the “[VSAN Membership for iSCSI](#)” procedure on page 4-11.

Detailed Steps

To assign VSAN membership for iSLB initiators, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator ip-address 10.1.1.3 switch(config-islb-init)#	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submenu.
Step 3	switch(config-islb-init)# vsan 3	Assigns the iSLB initiator node to a specified VSAN. Note You can assign this host to one or more VSANs.
	switch(config-islb-init)# no vsan 3	Removes the iSLB initiator from the specified VSAN.

**Note**

When an iSLB initiator is configured in any other VSAN (other than VSAN 1, the default VSAN), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

See the “[Configuring iSLB Using Device Manager](#)” procedure on page 4-56.

Configuring Metric for Load Balancing

You can assign a load metric to each initiator for weighted load balancing. The load calculated is based on the number of initiators on a given iSCSI interface. This feature accommodates initiators with different bandwidth requirements. For example, you could assign a higher load metric to a database server than to a web server. Weighted load balancing also accommodates initiators with different link speeds.

Also, you can configure initiator targets using the device alias or the pWWN. If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

For more information on load balancing, see the “[About Load Balancing Using VRRP](#)” section on page 4-19.

Choose **IP > iSCSI iSLB** in Device Manager and set the LoadMetric field to change the load balancing metric for an iSLB initiator.

See the “[Configuring iSLB Using Device Manager](#)” procedure on page 4-56.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To configure a weight for load balancing, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)#	Configures an iSLB initiator using the name of the initiator node and enters iSLB initiator configuration mode.
Step 3	switch(config-iscsi-init)# metric 100	Assigns 100 as the weight metric for this iSLB initiator.
Step 4	switch(config-iscsi-init)# no metric 100	Reverts to the default value (1000).

Configuring iSLB Initiator Targets

You can configure initiator targets using the device alias or the pWWN. You can also optionally specify one or more of the following optional parameters:

- Secondary pWWN
- Secondary device alias
- LUN mapping
- IQN
- VSAN identifier



Note The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.

In addition, you can disable auto-zoning.

If you configure an IQN for an initiator target, then that name is used to identify the initiator target. Otherwise, a unique IQN is generated for the initiator target.

Detailed Steps

To configure iSLB initiator targets, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator ip-address 10.1.1.3 switch(config-islb-init)#	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 3	<code>switch(config-iscsi-islb-init)# target pwwn 26:00:01:02:03:04:05:06</code>	Configures the iSLB initiator target using a pWWN with auto-zoning enabled (default).
	<code>switch(config-iscsi-islb-init)# target pwwn 26:00:01:02:03:04:05:06 no-zone</code>	Configures the iSLB initiator target using a pWWN with auto-zoning disabled.
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias</code>	Configures the iSLB initiator target using a device alias with auto-zoning enabled (default).
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias fc-lun 0x1234 iscsi-lun 0x2345</code>	Configures the iSLB initiator target using a device alias and optional LUN mapping. Note The CLI interprets the LUN identifier value as a hexadecimal value whether or not the <code>0x</code> prefix is included.
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias iqn-name iqn.1987-01.com.cisco.initiator</code>	Configures the iSLB initiator target using a device alias and an optional IQN.
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias sec-device-alias SecondaryAlias</code>	Configures the iSLB initiator target using a device alias and an optional secondary device alias.
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias sec-pwwn 26:01:02:03:04:05:06:07</code>	Configures the iSLB initiator target using a device alias and an optional secondary pWWN.
	<code>switch(config-iscsi-islb-init)# target device-alias SampleAlias vsan 10</code>	Configures the iSLB initiator target using a device alias and the VSAN identifier. Note The VSAN identifier is optional if the target is online. If the target is not online, the VSAN identifier is required.
	<code>switch(config-iscsi-init)# no target pwwn 26:00:01:02:03:04:05:06</code>	Removes the iSLB initiator target.

To configure additional iSLB initiator targets using Device Manager, follow these steps:

-
- Step 1 Choose **IP > iSCSI iSLB**.
You see the iSCSI iSLB dialog box.
 - Step 2 Click on the initiator you want to add targets to and click **Edit Initiator Specific Targets**.
You see the Initiator Specific Target dialog box.
 - Step 3 Click **Create** to create a new initiator target.
You see the Create Initiator Specific Target dialog box.
 - Step 4 Fill in the pWWN field with the initiator target pWWN.
 - Step 5 (Optional) Set the Name field to a globally unique identifier (IQN).
 - Step 6 (Optional) Check the **NoAutoZoneCreation** check box to disable auto-zoning.
 - Step 7 (Optional) Check the **TresspassMode** check box. See the “LUN Trespass for Storage Port Failover” section on page 4-25.
 - Step 8 (Optional) Check the **RevertToPrimary** check box to revert back to the primary port after an HA failover when the primary port comes back up.
 - Step 9 Set the PrimaryVsan to the VSAN for the iSLB initiator target.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 10** Click **Create** to create this iSLB initiator target.
- Step 11** If CFS is enabled, select **commit** from the CFS drop-down menu.

Configuring and Activating Zones for iSLB Initiators and Initiator Targets

You can configure a zone name where the iSLB initiators and initiator targets are added. If you do not specify a zone name, the IPS manager creates one dynamically.

iSLB zone sets have the following restrictions:

- Auto-zoning of the initiator with the initiator targets is enabled by default.
- A zone set must be active in a VSAN for auto-zones to be created in that VSAN.
- iSLB zone set activation might fail if another zone set activation is in process or if the zoning database is locked. Retry the iSLB zone set activation if a failure occurs. To avoid this problem, only perform only one zoning related operation (normal zones, IVR zones, or iSLB zones) at a time.
- Auto-zones are created when the zone set is activated and there has been at least one change in the zoneset. The activation has no effect if only the auto-zones have changed.



Caution

If IVR and iSLB are enabled in the same fabric, at least one switch in the fabric must have both features enabled. Any zoning related configuration or activation operation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, traffic might be disrupted in the fabric.

To configure the iSLB initiator optional auto-zone name and activate the zone set, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# islb initiator ip-address 10.1.1.3 switch(config-islb-init)#	Configures an iSLB initiator using its IPv4 address and enters iSLB initiator configuration submenu.
Step 3	switch(config-islb-init)# zonename IslbZone	Specifies the zone name where the initiators and the initiator targets are added (optional).
	switch(config-islb-init)# no zonename IslbZone	Removes the initiators and initiator targets from the zone and adds them to a dynamically created zone (default).
Step 4	switch(config-islb-init)# exit	Returns to configuration mode.
Step 5	switch(config)# islb zoneset activate	Activates zoning for the iSLB initiators and initiator targets with zoning enabled and creates auto-zones if no zone names are configured.
		Note This step is not required if CFS is enabled. CFS automatically activates the zone when the configuration changes are committed.

Choose **IP > iSCSI iSLB** in Device Manager and set the `autoZoneName` field to change the auto zone name for an iSLB initiator.

See the “Configuring iSLB Using Device Manager” procedure on page 4-56.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Examples

The following example shows the **show zoneset active** command output when the dynamically generated zone name is used.

```
switch# show zoneset active
zoneset name zoneset-1 vsan 1
  zone name ips_zone_5d9603bcff68008a6fc5862a6670ca09 vsan 1
    * fcid 0x010009 [ip-address 10.1.1.3]
      pwwn 22:00:00:04:cf:75:28:4d
      pwwn 22:00:00:04:cf:75:ed:53
      pwwn 22:00:00:04:cf:75:21:d5
      pwwn 22:00:00:04:cf:75:ee:59
    ...
```

The following example shows the **show zoneset active** command output when the configured zone name `IslbZone` is used.

```
switch# show zoneset active
zoneset name zoneset-1 vsan 1
  zone name ips_zone_IslbZone vsan 1
    ip-address 10.1.1.3
    pwwn 22:00:00:04:cf:75:28:4d
    pwwn 22:00:00:04:cf:75:ed:53
    pwwn 22:00:00:04:cf:75:21:d5
    pwwn 22:00:00:04:cf:75:ee:59
  ...
```

Restricting iSLB Initiator Authentication

By default, the iSLB initiator can use any user name in the RADIUS or local AAA database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSLB initiator name). The IPS module or MPS-14/2 module allows the initiator to log in as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password have been compromised.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator name iqn.1987-02.com.cisco.init switch(config-islb-init)#	Configures an iSLB initiator using the IQN of the initiator node and enters iSLB initiator configuration mode.
Step 3	switch(config-islb-init)# username user1	Restricts the initiator <code>iqn.1987-02.com.cisco.init</code> to only authenticate using <code>user1</code> as its CHAP user name. Tip Be sure to define <code>user1</code> as an iSCSI user in the local AAA database or the RADIUS server.

Choose **IP > iSCSI iSLB** in Device Manager and set the `AuthName` field to restrict an initiator to use a specific user name for CHAP authentication.

See the “Configuring iSLB Using Device Manager” procedure on page 4-56.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Mutual CHAP Authentication

In addition to the IPS module and MPS-14/2 module authentication of the iSLB initiator, the IPS module and MPS-14/2 module also support a mechanism for the iSLB initiator to authenticate the Cisco MDS switch's initiator target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSLB initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

Detailed Steps

To configure a per-initiator user name and password used by the switch to authenticate itself to an initiator, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb initiator name iqn.1987-02.com.cisco.initiator switch(config-islb-init)#	Configures an iSLB initiator using the name of the initiator node and enters iSLB initiator configuration mode.
Step 3	switch(config-islb-init)# mutual-chap username testuser password dcba12LKJ	Configures the switch user account (testuser) along with a password (dcba12LKJ) specified in clear text (default). The password is limited to 128 characters.
	switch(config-islb-init)# mutual-chap username testuser password 7 !@*asdfsdfjh!@df	Configures the switch user account (testuser) along with the encrypted password specified by 7 (!@*asdfsdfjh!@df).
Step 4	switch(config-iscsi-init)# no mutual-chap username testuser	Removes the switch authentication configuration.

Choose **IP > iSCSI iSLB** in Device Manager and set the Target Username and Target Password fields to configure a per-initiator user name and password used by the switch to authenticate itself to an initiator. See the “Configuring iSLB Using Device Manager” procedure on page 4-56.

Examples

Use the **show running-config** and the **show iscsi global** (see Example 4-6) commands to display the global configuration. Use the **show running-config** and the **show islb initiator configured** (see Example 4-14) commands to display the initiator specific configuration.

To verify the iSLB user name and mutual CHAP configuration, use the **show islb initiator configured** command.

```
switch# show islb initiator configured
iSCSI Node name is 10.1.1.3
Member of vsans: 3
User Name for login authentication: user1
User Name for Mutual CHAP: testuser
Load Balance Metric: 1000 Number of Initiator Targets: 1
Number of Initiator Targets: 1

Initiator Target: iqn.1987-05.com.cisco:05.ips-hac4
Port WWN 50:06:04:82:ca:e1:26:8d
Zoning Enabled
No. of LU mapping: 3
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
iSCSI LUN: 0x0001, FC LUN: 0x0001
iSCSI LUN: 0x0002, FC LUN: 0x0002
iSCSI LUN: 0x0003, FC LUN: 0x0003
```

Configuring Load Balancing Using VRRP

You must first configure VRRP on the Gigabit Ethernet interfaces on the switch that connect to the IP network before configuring VRRP for iSLB.

Detailed Steps

To configure VRRP load balancing using Device Manager, follow these steps:

-
- Step 1** Choose **IP > iSCSI iSLB**.
You see the iSCSI iSLB dialog box.
 - Step 2** Click the **VRRP** tab.
 - Step 3** Click **Create** to configure VRRP load balancing for iSLB initiators.
You see the Create iSCSI iSLB VRRP dialog box.
 - Step 4** Set the Vrid to the VRRP group number.
 - Step 5** Select either **ipv4** or **ipv6** and check the **LoadBalance** check box.
 - Step 6** Click **Create** to enable load balancing.
 - Step 7** If CFS is enabled, select **commit** from the CFS drop-down menu.
-

Enabling VRRP for Load Balancing

Detailed Steps

To enable or disable VRRP for iSLB, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# islb vrrp 10 load-balance	Enables iSLB VRRP for IPv4 VR group 10.
Step 3	switch(config)# no islb vrrp 10 load-balance	Disables iSLB VRRP for IPv4 VR group 10.
Step 4	switch(config)# islb vrrp ipv6 20 load-balance	Enables iSLB VRRP for IPv6 VR group 20.
Step 5	switch(config)# no islb vrrp ipv6 20 load-balance	Disables iSLB VRRP for IPv6 VR group 20.

Examples

To verify the iSLB VRRP load balancing configuration for IPv4, use the **show vrrp vr** command.

```
switch# show vrrp vr 1
      Interface VR IpVersion Pri   Time Pre State   VR IP addr
-----
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
GigE1/5 1 IPv4 100 1 s master 10.10.10.1
GigE1/6 1 IPv4 100 1 s master 10.10.10.1
```

To verify the iSLB VRRP load balancing configuration for IPv6, use the **show vrrp ipv6 vr** command.

```
switch# show vrrp ipv6 vr 1
Interface VR IpVersion Pri Time Pre State VR IP addr
-----
GigE6/2 1 IPv6 100 100cs master 5000:1::100
PortCh 4 1 IPv6 100 100cs master 5000:1::100
```

Distributing the iSLB Configuration Using CFS

This section contains the following:

- [Enabling iSLB Configuration Distribution, page 4-67](#)
- [Committing Changes to the Fabric, page 4-68](#)
- [Discarding Pending Changes, page 4-68](#)
- [Clearing a Fabric Lock, page 4-69](#)
- [Creating a Static iSCSI Virtual Target, page 4-69](#)
- [Enabling the Trespass Feature for a Static iSCSI, page 4-71](#)

Enabling iSLB Configuration Distribution

Detailed Steps

To enable CFS distribution of the iSLB configuration, follow these steps:

	Command	Purpose
Step 1	switch# confi g t	Enters configuration mode.
Step 2	switch(config)# isl b distribute	Enables iSLB configuration distribution.
	switch(config)# no isl b distribute	Disables (default) iSLB configuration distribution.

To enable CFS distribution of the iSLB configuration using Device Manager, follow these steps:

-
- Step 1** Choose **Admin > CFS**.
You see the CFS dialog box.
- Step 2** Set the Command field to **enable** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Committing Changes to the Fabric

To apply the pending iSLB configuration changes to the active configuration and to other MDS switches in the fabric, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the MDS switches throughout the fabric, the automatic zones are activated, and the fabric lock is released.

Detailed Steps

To commit iSLB configuration changes to other MDS switches in the fabric, activate iSLB automatic zones, and release the fabric lock, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# islb commit	Commits the iSLB configuration distribution, activates iSLB automatic zones, and releases the fabric lock.

To commit iSLB configuration changes to other MDS switches in the fabric, activate iSLB automatic zones, and release the fabric lock using Device Manager, follow these steps:

-
- Step 1 Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
 - Step 2 Set the Command field to **commit** for the iSLB feature.
 - Step 3 Click **Apply** to save this change.
-

Discarding Pending Changes

At any time, you can discard the pending changes to the iSLB configuration and release the fabric lock. This action has no affect on the active configuration on any switch in the fabric.

Detailed Steps

To discard the pending iSLB configuration changes and release the fabric lock, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# islb abort	Commits the iSLB configuration distribution.

To discard the pending iSLB configuration changes and release the fabric lock using Device Manager, follow these steps:

-
- Step 1 Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
 - Step 2 Set the Command field to **abort** for the iSLB feature.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click **Apply** to save this change.

Clearing a Fabric Lock

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

Restrictions

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

Detailed Steps

To release a fabric lock, issue the **clear islb session** command in EXEC mode using a login ID that has administrative privileges.

```
switch# clear islb session
```

To release a fabric lock using Device Manager, follow these steps:

- Step 1** Choose **Admin > CFS**.
You see the CFS Configuration dialog box.
- Step 2** Set the Command field to **clear** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
-

Creating a Static iSCSI Virtual Target

Detailed Steps

To create a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 3	<code>switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06</code>	Configures the primary port for this virtual target.
	<code>switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 secondary-pwwn 26:00:01:02:03:10:11:12</code>	Configures the primary and secondary ports for this virtual target.
	<code>switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 fc-lun 0x1 iscsi-lun 0x0 sec-lun 0x3</code>	Configures the primary port for this virtual target with LUN mapping and different LUN on the secondary Fibre Channel port. Note The CLI interprets the LUN identifier value as a hexadecimal value whether or not the 0x prefix is included.
	<code>switch(config-iscsi-tgt)# no pwwn 26:00:01:02:03:04:05:06</code>	Removes the primary port, secondary port, and LUN mapping configuration for this virtual target.
Step 4	<code>switch(config-iscsi-tgt)# revert-primary-port</code>	Configures the session failover redundancy for this virtual-target to switch all sessions back to primary port when the primary port comes back up.
Step 5	<code>switch(config-iscsi-tgt)# no revert-primary-port</code>	Directs the switch to continue using the secondary port for existing sessions and to use the primary port for new sessions (default).

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

-
- Step 1** Click **IP > iSCSI**.
You see the iSCSI configuration.
 - Step 2** Click the **Targets** tab to display a list of existing iSCSI targets shown.
 - Step 3** Click **Create** to create an iSCSI target.
You see the Create iSCSI Targets dialog box.
 - Step 4** Set the iSCSI target node name in the iSCSI Name field, in IQN format.
 - Step 5** Set the Port WWN field for the Fibre Channel target port you are mapping.
 - Step 6** Click the **Select from List** radio button and set the iSCSI initiator node names or IP addresses that you want this virtual iSCSI target to access, or click the **All** radio button to let the iSCSI target access all iSCSI initiators. See the “iSCSI Access Control” section on page 4-11.
 - Step 7** Click the **Select from List** radio button and check each interface you want to advertise the iSCSI targets on or choose the **All** radio button to advertise all interfaces.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 8 [Click Apply to save this change.](#)

Enabling the Trespass Feature for a Static iSCSI

Detailed Steps

To enable the trespass feature for a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator switch(config-iscsi-tgt)#	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
Step 3	switch(config-iscsi-tgt)# pwwn 50:00:00:a1:94:cc secondary-pwwn 50:00:00:a1:97:ac	Maps a virtual target node to a Fibre Channel target and configures a secondary pWWN.
Step 4	switch(config-iscsi-tgt)# trespass	Enables the trespass feature.
	switch(config-iscsi-tgt)# no trespass	Disables the trespass feature (default).

In Device Manager, choose **IP > iSCSI**, select the **Targets** tab, and check the **Trespass Mode** check box to enable the trespass feature for a static iSCSI virtual target.

```
switch# show iscsi virtual-target iqn.1987-02.com.cisco.initiator
target: 1987-02.com.cisco.initiator
  Port WWN 10:20:10:00:56:00:70:50
  Configured node
  all initiator permit is disabled
  trespass support is enabled
```

Configuring iSCSI Authentication

This section provides configuration information on iSCSI authentication. It includes the following authentication procedures:

- [Configuring No Authentication, page 4-72](#)
- [Configuring CHAP with Local Password Database, page 4-72](#)
- [Configuring CHAP with External RADIUS Server, page 4-73](#)



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 entering any command.



Caution

Changing the authentication of an iSCSI interface that is part of an iSLB VRRP group impacts load balancing on the interface. See the [“Changing iSCSI Interface Parameters and the Impact on Load Balancing”](#) section on page 4-21.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring No Authentication

To configure a network with no authentication, set the iSCSI authentication method to **none**

Detailed Steps

In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane. Select the **Globals** tab and set the AuthMethod drop-down menu to **none** and click **Apply Changes**.

```
switch(config)# iscsi authentication none
```

Configuring CHAP with Local Password Database

Detailed Steps

To configure authentication using the CHAP option with the local password database, follow these steps:

Step 1 Set the AAA authentication to use the local password database for the iSCSI protocol.

```
switch(config)# aaa authentication iscsi default local
```

Step 2 Set the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

Step 3 Configure the user names and passwords for iSCSI users.

```
switch(config)# username iscsi-user password abcd iscsi
```



Note If you do not specify the **iscsi** option, the user name is assumed to be a Cisco MDS switch user instead of an iSCSI user.

Step 4 Verify the global iSCSI authentication setup.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----Verify
  Import FC Target: Disabled
  ...
```

To configure authentication using the CHAP option with the local password database, follow these steps:

Step 1 Set the AAA authentication to use the local password database for the iSCSI protocol:

- a. In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
- b. Click the **Applications** tab in the Information pane.
- c. Check the **Local** check box for the iSCSI row and click **Apply Changes**

Step 2 Set the iSCSI authentication method to require CHAP for all iSCSI clients:

- a. In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
- b. Click the **Globals** tab in the Information pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- c. Set the AuthMethod drop-down menu to **chap** and click **Apply Changes**.
- Step 3** Configure the user names and passwords for iSCSI users:
- a. In Device Manager, choose **Security > iSCSI**.
 - b. Set the Username, Password and Confirm Password fields.
 - c. Click **Create** to save these changes.
- Step 4** Verify the global iSCSI authentication setup:
- a. In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - b. Click the **Globals** tab in the Information pane.
-

Configuring CHAP with External RADIUS Server

Detailed Steps

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

- Step 1** Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server

```
switch(config)# radius-server key mds-1
```

- Step 2** Configure the RADIUS server IP address by performing one of the following:

- Configure an IPv4 address.

```
switch(config)# radius-server host 10.1.1.10
```

- Configure an IPv6 address.

```
switch(config)# radius-server host 2001:0DB8:800:200C::417A
```

- Step 3** Configure the RADIUS server group IP address by performing one of the following:

- Configure an IPv4 address.

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 10.1.1.1
```

- Configure an IPv6 address.

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 001:0DB8:800:200C::4180
```

```
switch(config)# aaa authentication iscsi default group iscsi-radius-group
```

- Step 4** Set up the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 5** Verify that the global iSCSI authentication setup is for CHAP.

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP          <----- Verify CHAP
  ....
```

- Step 6** Verify that the AAA authentication information is for iSCSI.

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch# show aaa authentication
      default: local
      console: local
      iscsi: group iscsi-radius-group    <----- Group name
      dhchap: local

switch# show radius-server groups
total number of groups:2

following RADIUS server groups are configured:
  group radius:
      server: all configured radius servers
  group iscsi-radius-group:
      server: 10.1.1.1 on auth-port 1812, acct-port 1813

switch# show radius-server
Global RADIUS shared secret:mds-1    <----- Verify secret
....

following RADIUS servers are configured:
  10.1.1.1:                            <----- Verify the server IPv4 address
      available for authentication on port:1812
      available for accounting on port:1813
```

- Step 1** Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.
 - Click the **Default** tab in the Information pane.
 - Set the AuthKey field to the default password and click the **Apply Changes** icon.
- Step 2** Configure the RADIUS server IP address:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA > RADIUS** in the Physical Attributes pane.
 - Click the **Server** tab in the Information pane and click **Create Row**.
 - Set the Index field to a unique number.
 - Set the IP Type radio button to **ipv4** or **ipv6**.
 - Set the Name or IP Address field to the IP address of the RADIUS server and click **Create**.
- Step 3** Create a RADIUS server group and add the RADIUS server to the group:
- In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
 - Select the **Server Groups** tab in the Information pane and click **Create Row**.
 - Set the Index field to a unique number.
 - Set the Protocol radio button to **radius**.
 - Set the Name field to the server group name.
 - Set the ServerIDList to the index value of the RADIUS server (as created in Step 2 c.) and click **Create**.
- Step 4** Set up the authentication verification for the iSCSI protocol to go to the RADIUS server.
- In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
 - Click the **Applications** tab in the Information pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- c. Right-click on the iSCSI row in the Type, SubType, Function column.
 - d. Set the ServerGroup IDList to the index value of the Server Group (as created in Step 3 c) and click **Create**.
- Step 5** Set up the iSCSI authentication method to require CHAP for all iSCSI clients.
- a. In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - b. Select **chap** from the AuthMethod drop-down menu.
 - c. Click the **Apply Changes** icon.
- Step 6** In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
- Step 7** Click the **Globals** tab in the Information pane to verify that the global iSCSI authentication setup is for CHAP.
- Step 8** In Cisco DCNM-SAN, choose **Switches > Security > AAA** in the Physical Attributes pane.
- Step 9** Click the **Applications** tab in the Information pane to verify the AAA authentication information for iSCSI.

To configure an iSCSI RADIUS server, follow these steps:

- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
 - Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
 - Step 3** Configure the iSCSI users and passwords on the RADIUS server.
-

Creating an iSNS Client Profile

Detailed Steps

To create an iSNS profile, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# isns profile name MyIsns switch(config-isns-profile)#	Creates a profile called MyIsns.
Step 3	switch(config-isns-profile)# server 10.10.100.211	Specifies an iSNS server IPv4 address for this profile.
Step 4	switch(config-isns-profile)# no server 10.10.100.211	Removes a configured iSNS server from this profile.
Step 5	switch(config-isns-profile)# server 2003::11	Specifies an iSNS server IPv6 address for this profile.
Step 6	switch(config-isns-profile)# no server 10.20.100.211	Removes a configured iSNS server from this profile.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To create an iSNS profile, follow these steps:

-
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
You see the iSCSI configuration in the Information pane.
 - Step 2** Select the **iSNS** tab.
 - Step 3** You see the iSNS profiles configured.
 - Step 4** Click the **Create Row** icon.
You see the Create iSNS Profiles dialog box.
 - Step 5** Set the ProfileName field to the iSNS profile name that you want to create.
 - Step 6** Set the ProfileAddr field to the IP address of the iSNS server.
 - Step 7** Click **Create** to save these changes.
-

To remove an iSNS profile, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# no isns profile name OldIsns	Removes a configured iSNS profile called OldIsns.

To delete an iSNS profile, follow these steps:

-
- Step 1** Choose **End Devices > iSCSI** from the Physical Attributes pane.
You see the iSCSI configuration in the Information pane.
 - Step 2** Select the **iSNS** tab.
You see the iSNS profiles configured.
 - Step 3** Right-click the profile that you want to delete and click the **Delete Row** icon.
-

To tag a profile to an interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 4/1 switch(config-if)#	Configures the specified Gigabit Ethernet interface.
Step 3	switch(config-if)# isns MyIsns	Tags a profile to an interface.

To tag a profile to an interface, follow these steps:

-
- Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
You see the Gigabit Ethernet configuration in the Information pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Click the **iSNS** tab.
You see the iSNS profiles configured for these interfaces.
- Step 3** Set the iSNS ProfileName field to the iSNS profile name that you want to add to this interface.
- Step 4** Click the **Apply Changes** icon to save these changes.

To untag a profile from an interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 5/1 switch(config-if)#	Configures the specified Gigabit Ethernet interface.
Step 3	switch(config-if)# no isns OldIsns	Untags a profile from an interface.

Use the **isns reregister** command in EXEC mode to reregister associated iSNS objects with the iSNS server.

```
switch# isns reregister gigabitethernet 1/4
switch# isns reregister port-channel 1
```

To untag a profile from an interface using Cisco DCNM-SAN, follow these steps:

- Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** Click the **iSNS** tab.
You see the iSNS profiles configured for these interfaces.
- Step 3** Right-click the iSNS ProfileName field that you want to untag and delete the text in that field.
- Step 4** Click the **Apply Changes** icon to save these changes.

Configuring iSNS Servers

This section describe how to configure an iSNS server on a Cisco MDS 9000 Family switch.

This section includes the following topics:

- [Enabling the iSNS Server, page 4-78](#)
- [iSNS Configuration Distribution, page 4-78](#)
- [Configuring the ESI Retry Count, page 4-79](#)
- [Configuring the Registration Period, page 4-80](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Enabling the iSNS Server

Before the iSNS server feature can be enabled, iSCSI must be enabled (see the “[Enabling iSCSI](#)” section on page 4-35). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

Detailed Steps

To enable the iSNS server, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# isns-server enable	Enables the iSNS server.
	switch(config)# no isns-server enable	Disables (default) the iSNS server.

To enable the iSNS server, follow these steps:

-
- Step 1 Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
 - Step 2 Click the **Control** tab and select **enable** from the Command drop-down menu for the iSNS server feature.
 - Step 3 Click the **Apply Changes** icon to save this change.
-



Note

If you are using VRRP IPv4 addresses for discovering targets from iSNS clients, ensure that the IP address is created using the **secondary** option (see “[Adding Virtual Router IP Addresses](#)” section on page 5-18).

iSNS Configuration Distribution

You can use the CFS infrastructure to distribute the iSCSI initiator configuration to iSNS servers across the fabric. This allows the iSNS server running on any switch to provide a querying iSNS client a list of iSCSI devices available anywhere on the fabric. For information on CFS, see the *System Management Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS System Management Configuration Guide.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To enable iSNS configuration distribution using, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# isns distribute	Uses the CFS infrastructure to distribute the iSCSI virtual target configuration to all switches in the fabric.
	switch(config)# no isns distribute	Stops (default) the distribution of iSCSI virtual target configuration to all switches in the fabric.

To enable iSNS configuration distribution, follow these steps:

-
- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **CFS** tab and select **enable** from the Admin drop-down menu for iSNS.
- Step 3** Select **enable** from the Global drop-down menu for iSNS.
- Step 4** Click the **Apply Changes** icon to save this change.
-

Configuring the ESI Retry Count

The iSNS client registers information with its configured iSNS server using an iSNS profile. At registration, the client can indicate an entity status inquiry (ESI) interval of 60 seconds or more. If the client registers with an ESI interval set to zero (0), then the server does not monitor the client using ESI. In such cases, the client's registrations remain valid until explicitly deregistered or the iSNS server feature is disabled.

The ESI retry count is the number of times the iSNS server queries iSNS clients for their entity status. The default ESI retry count is 3. The client sends the server a response to indicate that it is still alive. If the client fails to respond after the configured number of retries, the client is deregistered from the server.

Detailed Steps

To configure the ESI retry count for an iSNS server, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# isns esi retries 6	Configures the ESI to retry contacting the client up to 6 times. The range is 1 to 10.
	switch(config)# no isns esi retries 6	Reverts to the default value of 3 retries.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring the Registration Period

The iSNS client specifies the registration period with the iSNS Server. The iSNS Server keeps the registration active until the end of this period. If there are no commands from the iSNS client during this period, then the iSNS server removes the client registration from its database.

If the iSNS client does not specify a registration period, the iSNS server assumes a default value of 0, which keeps the registration active indefinitely. You can also manually configure the registration period on the MDS iSNS Server.

Detailed Steps

To configure the registration period on an iSNS Server, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# isns registration period 300	Configures the registration to be active for 300 seconds. The permissible registration period is between 0 to 65536 seconds.
	switch(config)# no isns registration period	Reverts to the client registered timeout value, or the default value of 0.

To configure the registration period on an iSNS Server, follow these steps:

-
- Step 1** Choose [End Devices > iSNS](#).
- You see the iSNS configuration in the Information pane.
- Step 2** Click the [Servers](#) tab.
- You see the configured iSNS servers.
- Step 3** Set the [ESI NonResponse Threshold](#) field to the ESI retry count value.
- Step 4** Click the [Apply Changes](#) icon to save this change.
-

Configuring iSNS Cloud Discovery

This section describes how to configure iSNS cloud discovery and includes the following topics:

- [Enabling iSNS Cloud Discovery](#), page 4-81
- [Initiating On-Demand iSNS Cloud Discovery](#), page 4-81
- [Configuring Automatic iSNS Cloud Discovery](#), page 4-82
- [Configuring iSNS Cloud Discovery Message Types](#), page 4-83

Send documentation comments to dcnm-san-docfeedback@cisco.com

Enabling iSNS Cloud Discovery

Detailed Steps

To enable iSNS cloud discovery, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cloud-discovery enable	Enables iSNS cloud discovery.
	switch(config)# no cloud-discovery enable	Disables (default) iSNS cloud discovery.

To enable iSNS cloud discovery, follow these steps:

-
- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **Control** tab and select **enable** from the Command drop-down menu for the cloud discovery feature.
- Step 3** Click the **Apply Changes** icon to save this change.
-

Initiating On-Demand iSNS Cloud Discovery

Detailed Steps

To initiate on-demand iSNS cloud discovery, use the **cloud discover** command in EXEC mode.

The following example shows how to initiate on-demand cloud discovery for the entire fabric:

```
switch# cloud discover
```

To initiate on-demand iSNS cloud discovery, follow these steps:

-
- Step 1** Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
- Step 2** Click the **Cloud Discovery** tab and check the **Manual Discovery** check box.
- Step 3** Click the **Apply Changes** icon to save this change.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Automatic iSNS Cloud Discovery

Detailed Steps

To configure automatic iSNS cloud discovery, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cloud discovery auto	Enables (default) automatic iSNS cloud discovery.
	switch(config)# no cloud discovery auto	Disables automatic iSNS cloud discovery.

To configure automatic iSNS cloud discovery, follow these steps:

-
- Step 1 Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
 - Step 2 Click the **Cloud Discovery** tab and check the **AutoDiscovery** check box.
 - Step 3 Click the **Apply Changes** icon to save this change.
-

Configuring iSNS Cloud Discovery Distribution

Detailed Steps

To configure iSNS cloud discovery distribution using CFS, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cloud discovery fabric distribute	Enables (default) iSNS cloud discovery fabric distribution.
	switch(config)# no cloud discovery fabric distribute	Disables iSNS cloud discovery fabric distribution.

To configure iSNS cloud discovery CFS distribution, follow these steps:

-
- Step 1 Choose **End Devices > iSNS**.
You see the iSNS configuration in the Information pane.
 - Step 2 Click the **CFS** tab and select **enable** from the Admin drop-down menu for the cloud discovery feature.
 - Step 3 Select **enable** from the Global drop-down menu for the cloud discovery feature.
 - Step 4 Click the **Apply Changes** icon to save this change.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring iSNS Cloud Discovery Message Types

You can configure iSNS cloud discovery the type of message to use. By default, iSNS cloud discovery uses ICMP.

To configure iSNS cloud discovery message types, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cloud discovery message icmp	Enables (default) iSNS cloud discovery using ICMP messages. Note Only ICMP messages are supported.

Verifying iSCSI Configuration

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

Command	Purpose
show interface iscsi 4/1	Displays the iSCSI Interface Information
show iscsi stats iscsi 2/1	Display brief iSCSI statistics for an iSCSI interface.
show iscsi stats iscsi 2/1 detail	Displays detailed iSCSI statistics for the iSCSI interface.
show interface iscsi 4/1	Displays proxy initiator information for the iSCSI interface with system-assigned WWNs.
show interface iscsi 4/2	Displays proxy initiator information for the iSCSI interface with user-assigned WWNs.
show iscsi global	Displays the current global iSCSI configuration and date.
show iscsi session.	Displays brief information of all iSCSI sessions.
show iscsi session initiator 10.10.100.199 target VT1	Displays brief information about the specified iSCSI session.
show iscsi session initiator 10.10.100.199 target VT1 detail	Displays detailed information about the specified iSCSI session.
switch# show iscsi initiator	Displays information about connected iSCSI initiators.
show iscsi initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail	Displays detailed information about the iSCSI initiator.
show fcns database	Displays the FCNS database contents.
show fcns database detail.	Displays the FCNS database in detail.
show iscsi initiator configured	Displays information about configured initiators.
show iscsi virtual-target	Displays exported targets.
show user-account iscsi	Displays iSCSI user names.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Command	Purpose
<code>show islb initiator configured</code>	Verifies the iSLB initiator configuration.
<code>show islb initiator configured</code>	Verifies the iSLB target configuration.
<code>show zoneset active</code>	Shows the <code>show zoneset active</code> command output when the dynamically generated zone name is used.
<code>show zoneset active</code>	Shows the <code>show zoneset active</code> command output when the configured zone name <code>IslbZone</code> is used.
<code>show islb initiator configured</code>	Verifies the iSLB user name and mutual CHAP configuration.
<code>show vrrp vr 1</code>	Verifies the iSLB VRRP load balancing configuration for IPv4.
<code>show vrrp ipv6 vr 1</code>	Verifies the iSLB VRRP load balancing configuration for IPv6.
<code>show islb vrrp summary vr 30</code>	Displays VRRP load balancing information.
<code>show islb pending</code>	Displays the pending configuration changes.
<code>show islb pending-diff</code>	Displays the differences between the pending configuration and the current configuration.
<code>show islb status</code>	Displays the iSLB CFS status.
<code>show islb cfs-session status</code>	Displays the status of the iSLB CFS distribution session.
<code>show islb merge status</code>	Displays the iSLB CFS merge status.
<code>show isns profile</code>	Displays information for configured iSNS profiles.
<code>show isns profile ABC</code>	Displays a specified iSNS profile.
<code>show isns profile counters</code>	Displays configured profiles with iSNS statistics.
<code>show isns profile ABC counters</code>	Displays iSNS statistics for a specified profile.
<code>show isns query ABC gigabitethernet 2/3</code>	Displays iSNS queries.
<code>show interface gigabitethernet 2/3</code>	Displays tagged iSNS interfaces.
<code>show isns config</code>	Displays the iSNS server configuration of ESI interval and database contents.
<code>show isns database</code>	Displays explicitly registered objects.
<code>show isns database full</code>	Displays the full database with both registered and configured nodes and portal.
<code>show isns database virtual-targets local</code>	Displays the virtual target information in the local switch.
<code>show isns database virtual-targets switch 20:00:00:0d:ec:01:04:40</code>	Displays virtual target for a specified switch.
<code>show isns node all</code>	Displays explicitly registered objects.
<code>show isns node name iqn.com.cisco.disk1</code>	Displays the specified node.
<code>show isns node all detail</code>	Displays the attribute details for all nodes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Command	Purpose
<code>show isns portal all</code>	Displays the attribute information for all portals.
<code>show isns portal all detail</code>	Displays detailed attribute information for all portals.
<code>show isns portal virtual</code>	Displays virtual portals.
<code>show isns portal virtual switch 20:00:00:0d:ec:01:04:40</code>	Displays virtual portals for the specified switch.
<code>show isns portal virtual switch 20:00:00:0d:ec:01:04:40 detail</code>	Displays detailed information for the virtual portals in the specified switch.
<code>show isns entity</code>	Displays all registered entries.
<code>show isns entity all</code>	Displays all entities in the database.
<code>show isns entity id dp-204</code>	Displays the entity with the specified ID.
<code>show isns entity all detail</code>	Displays detailed information for all entities in the database.
<code>show isns entity virtual</code>	Displays the virtual entities.
<code>show isns iscsi global config switch 20:00:00:05:ec:01:04:00</code>	Displays the import target settings for the specified switch.
<code>show isns iscsi global config all</code>	Displays the import target settings for all switches.
<code>show cfs peers name isns</code>	Displays the CFS peer switch information for the iSNS application.
<code>show cloud discovery status</code>	Verifies the status of the cloud discovery operation.
<code>show cloud membership all</code>	Verifies the cloud membership for the switch.
<code>show cloud membership unresolved</code>	Verifies the unresolved membership on the switch.
<code>show cloud discovery statistics</code>	Displays the statistics for the cloud discovery operation.
<code>show cloud discovery config</code>	Shows cloud discovery config command.
<code>show islb vrrp summary</code>	Shows the initial load distribution for three initiators with the default load metric value.

Use the **show iscsi** command to obtain detailed information about iSCSI configurations.

This section includes the following topics:

- Displaying iSCSI Interfaces, page 4-86
- Displaying iSCSI Statistics, page 4-87
- Displaying Proxy Initiator Information, page 4-88
- Displaying Global iSCSI Information, page 4-90
- Displaying iSCSI Sessions, page 4-90
- Displaying iSCSI Initiators, page 4-91
- Displaying iSCSI Virtual Targets, page 4-95
- Displaying iSCSI User Information, page 4-95

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- Displaying iSLB VRRP Information, page 4-95
- Displaying Pending iSLB Configuration Changes, page 4-96
- Displaying iSLB CFS Status, page 4-96
- Displaying iSLB CFS Distribution Session Status, page 4-96
- Displaying iSLB CFS Merge Status, page 4-96
- Verifying iSNS Client Configuration, page 4-97
- Verifying the iSNS Server Configuration, page 4-98
- Verifying Automatic iSNS Cloud Discovery Configuration, page 4-104
- Verifying Cloud Discovery Status, page 4-105
- Verifying Cloud Discovery Membership, page 4-105
- Displaying Cloud Discovery Statistics, page 4-105

Displaying iSCSI Interfaces

Use the **show iscsi interface** command to view the summary, counter, description, and status of the iSCSI interface. Use the output to verify the administrative mode, the interface status, TCP parameters currently used, and brief statistics.

Example 4-1 Displays the iSCSI Interface Information

```
switch# show interface iscsi 4/1
iscsi4/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:cf:00:0c:85:90:3e:80
  Admin port mode is ISCSI
  Port mode is ISCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0 (discovery session: 0)
  Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is enabled
    QOS code point is 0
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 70000 kbps
    Estimated round trip time is 1000 usec
    Send buffer size is 4096 KB
    Congestion window monitoring is enabled, burst size is 50 KB
    Configured maximum jitter is 500 us
  Forwarding mode: store-and-forward
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : disabled
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 0 packets, 0 bytes
      Command 0 pdus, Data-out 0 pdus, 0 bytes
    Output 0 packets, 0 bytes
```


Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Response 0 pdus (with sense 0), R2T 0 pdus
Data-in 0 pdus, 0 bytes
```

Displaying iSCSI Statistics

Use the **show iscsi stats** command to view brief or detailed iSCSI statistics per iSCSI interface. See Example 4-2 and Example 4-3.

Example 4-2 displays iSCSI throughput on an IPS port in both inbound and outbound directions. It also displays the number of different types of iSCSI PDU received and transmitted by this IPS port.

Example 4-2 Display Brief iSCSI Statistics for an iSCSI Interface

```
switch# show iscsi stats iscsi 2/1
iscsi2/1
  5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
iSCSI statistics
  974756 packets input, 142671620 bytes
    Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
  bytes
  output 1022920 packets, 143446248 bytes
    Response 2352 pdus (with sense 266), R2T 1804 pdus
    Data-in 90453 pdus, 92458248 bytes
```

Example 4-3 displays detailed iSCSI statistics for an IPS port. Along with the traffic rate and the number of each iSCSI PDU type, it shows the number of FCP frames received and forwarded, the number of iSCSI login attempts, successes, and failures. It also shows the number of different types of iSCSI PDUs sent and received that are noncritical or occur less frequently, such as NOP in and out (NOP-In and NOP-Out), text request and response (Text-REQ and Text-RESP), and task management request and response (TMF-REQ and TMF-RESP).

Various types of errors and PDU or frame drop occurrences are also counted and displayed. For example, Bad header digest shows the number of iSCSI PDUs received that have a header digest that fails CRC verification. The iSCSI Drop section shows the number of PDUs that were dropped because of reasons such as target down, LUN mapping fail, Data CRC error, or unexpected Immediate or Unsolicited data. These statistics are helpful for debugging purposes when the feature is not working as expected.

The last section, Buffer Stats, gives statistics on the internal IPS packet buffer operation. This section is for debugging purposes only.

Example 4-3 Displays Detailed iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 2/1 detail
iscsi2/1
  5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
iSCSI statistics
  974454 packets input, 142656516 bytes
    Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
  bytes
  output 1022618 packets, 143431144 bytes
    Response 2352 pdus (with sense 266), R2T 1804 pdus
    Data-in 90453 pdus, 92458248 bytes
iSCSI Forward:
  Command:2352 PDUs (Rcvd:2352)
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

Data-Out (Write):16236 PDUs (Rcvd 44198), 0 fragments, 92364800 bytes, unsolicited 0
bytes
FCP Forward:
  Xfer_rdy:1804 (Rcvd:1804)
  Data-In:90453 (Rcvd:90463), 92458248 bytes
  Response:2352 (Rcvd:2362), with sense 266
  TMF Resp:0

iSCSI Stats:
  Login:attempt:13039, succeed:110, fail:12918, authen fail:0
  Rcvd:NOP-Out:914582, Sent:NOP-In:914582
    NOP-In:0, Sent:NOP-Out:0
    TMF-REQ:0, Sent:TMF-RESP:0
    Text-REQ:18, Sent:Text-RESP:27
  SNACK:0
  Unrecognized Opcode:0, Bad header digest:0
  Command in window but not next:0, exceed wait queue limit:0
  Received PDU in wrong phase:0
  SCSI Busy responses:0
  Immediate data failure::Separation:0
  Unsolicited data failure::Separation:0, Segment:0
    Add header:0
  Sequence ID allocation failure:0
FCP Stats:
  Total:Sent:47654
    Received:96625 (Error:0, Unknown:0)
  Sent:PLOGI:10, Rcvd:PLOGI_ACC:10, PLOGI_RJT:0
    PRLI:10, Rcvd:PRLI_ACC:10, PRLI_RJT:0, Error:0, From initiator:0
    LOGO:4, Rcvd:LOGO_ACC:0, LOGO_RJT:0
    PRLO:4, Rcvd:PRLO_ACC:0, PRLO_RJT:0
    ABTS:0, Rcvd:ABTS_ACC:0
    TMF REQ:0
    Self orig command:10, Rcvd:data:10, resp:10
  Rcvd:PLOGI:156, Sent:PLOGI_ACC:0, PLOGI_RJT:156
    LOGO:0, Sent:LOGO_ACC:0, LOGO_RJT:0
    PRLI:8, Sent:PRLI_ACC:8, PRLI_RJT:0
    PRLO:0, Sent:PRLO_ACC:0, PRLO_RJT:0
    ADISC:0, Sent:ADISC_ACC:0, ADISC_RJT:0
    ABTS:0

iSCSI Drop:
  Command:Target down 0, Task in progress 0, LUN map fail 0
    CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
    No task:0
  Data-Out:0, Data CRC Error:0
  TMF-Req:0, No task:0
  Unsolicited data:0, Immediate command PDU:0
FCP Drop:
  Xfer_rdy:0, Data-In:0, Response:0

Buffer Stats:
  Buffer less than header size:0, Partial:45231, Split:322
  Pullup give new buf:0, Out of contiguous buf:0, Unaligned m_data:0

```

Displaying Proxy Initiator Information

If the proxy initiator feature is enabled in the iSCSI interface, use the **show interface iscsi** command to display configured proxy initiator information (see Example 4-4 and Example 4-5).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 4-4 Displays Proxy Initiator Information for the iSCSI Interface with System-Assigned WWNs

```
switch# show interface iscsi 4/1
iscsi4/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:c1:00:05:30:00:a7:9e
  Admin port mode is ISCSI
  Port mode is ISCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0, Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is disabled
    QOS code point is 0
  Forwarding mode: pass-thru
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : enabled<-----Proxy initiator is enabled
    nWWN is 28:00:00:05:30:00:a7:a1 (system-assigned)<----System-assigned nWWN
    pWWN is 28:01:00:05:30:00:a7:a1 (system-assigned)<---- System-assigned pWWN
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 7 packets, 2912 bytes
      Command 0 pdus, Data-out 0 pdus, 0 bytes
    Output 7 packets, 336 bytes
      Response 0 pdus (with sense 0), R2T 0 pdus
      Data-in 0 pdus, 0 bytes
```

Example 4-5 Displays Proxy Initiator Information for the iSCSI Interface with User-Assigned WWNs

```
switch# show interface iscsi 4/2
iscsi4/2 is up
  Hardware is GigabitEthernet
  Port WWN is 20:c1:00:05:30:00:a7:9e
  Admin port mode is ISCSI
  Port mode is ISCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0, Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is disabled
    QOS code point is 0
  Forwarding mode: pass-thru
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : enabled
    nWWN is 11:11:11:11:11:11:11:11 (manually-configured)<----User-assigned nWWN
    pWWN is 22:22:22:22:22:22:22:22 (manually-configured)<----User-assigned pWWN
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 7 packets, 2912 bytes
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Command 0 pdus, Data-out 0 pdus, 0 bytes
Output 7 packets, 336 bytes
Response 0 pdus (with sense 0), R2T 0 pdus
Data-in 0 pdus, 0 bytes
```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See Example 4-6.

Example 4-6 Displays the Current Global iSCSI Configuration and State

```
switch# show iscsi global
iSCSI Global information
Authentication: CHAP, NONE
Import FC Target: Enabled
Initiator idle timeout: 300 seconds
Number of target node: 0
Number of portals: 11
Number of session: 0
Failed session: 0, Last failed initiator name:
```

Displaying iSCSI Sessions

Use the **show iscsi session** command to view details about the current iSCSI sessions in the switch. Without parameters, this command displays all sessions. The output can be filtered by specifying an initiator, a target, or both.

Example 4-7 displays one iSCSI initiator configured based on the IQN (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on its IPv4 address (10.10.100.199).

Example 4-7 Displays Brief Information of All iSCSI Sessions

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
Initiator ip addr (s): 10.10.100.116
Session #1
Discovery session, ISID 00023d000043, Status active

Session #2
Target VT1
VSAN 1, ISID 00023d000046, Status active, no reservation

Session #3
Target VT2
VSAN 1, ISID 00023d000048, Status active, no reservation

Initiator 10.10.100.199
Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
Session #1
Target VT2
VSAN 1, ISID 246700000000, Status active, no reservation

Session #2
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Target VT1
VSAN 1, ISID 246b00000000, Status active, no reservation
```

```
Session #3
Target iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
VSAN 1, ISID 246e00000000, Status active, no reservation
```

Example 4-8 and Example 4-9 display the iSCSI initiator configured based on its IPv4 address (10.10.100.199).

Example 4-8 Displays Brief Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1
Initiator 10.10.100.199
Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
Session #1
Target VT1
VSAN 1, ISID 246b00000000, Status active, no reservation
```

Example 4-9 Displays Detailed Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1 detail
Initiator 10.10.100.199 (oasis-qa)
Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
Session #1 (index 3)
Target VT1
VSAN 1, ISID 246b00000000, TSIH 384, Status active, no reservation
Type Normal, ExpCmdSN 39, MaxCmdSN 54, Barrier 0
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
Stats:
PDU: Command: 38, Response: 38
Bytes: TX: 8712, RX: 0
Number of connection: 1
Connection #1
Local IP address: 10.10.100.200, Peer IP address: 10.10.100.199
CID 0, State: LOGGED_IN
StatSN 62, ExpStatSN 0
MaxRecvDSLength 1024, our_MaxRecvDSLength 1392
CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
Version Min: 2, Max: 2
FC target: Up, Reorder PDU: No, Marker send: No (int 0)
Received MaxRecvDSLen key: No
```

Displaying iSCSI Initiators

Use the **show iscsi initiator** command to display information about all initiators connected to an iSCSI interface in the switch. The information can be filtered to display only the desired iSCSI initiator by specifying the initiator name. Detailed output of the iSCSI initiator can be obtained by specifying the **detail** option. The **iscsi-session** (and optionally **detail**) parameter displays only iSCSI session information. The **fcp-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See Example 4-10 and Example 4-11.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 4-10 Displays Information About Connected iSCSI Initiators

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1
  Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag: 0x180
    VSAN ID 1, FCID 0x6c0202
    VSAN ID 2, FCID 0x6e0000
    VSAN ID 10, FCID 0x790000

iSCSI Node name is 10.10.100.199
  iSCSI Initiator name: iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  iSCSI alias name: oasis-qa
  Node WWN is 22:03:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 5
  Number of Virtual n_ports: 1
  Virtual Port WWN is 22:00:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag: 0x180
    VSAN ID 5, FCID 0x640000
    VSAN ID 1, FCID 0x6c0203
```

Example 4-11 Displays Detailed Information About the iSCSI Initiator

```
switch# show iscsi initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1

Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
  Interface iSCSI 4/1, Portal group tag is 0x180
  VSAN ID 1, FCID 0x6c0202
  1 FC sessions, 1 iSCSI sessions
  iSCSI session details <-----iSCSI session details
  Target: VT1
  Statistics:
    PDU: Command: 0, Response: 0
    Bytes: TX: 0, RX: 0
    Number of connection: 1
  TCP parameters
    Local 10.10.100.200:3260, Remote 10.10.100.116:4190
    Path MTU: 1500 bytes
    Retransmission timeout: 310 ms
    Round trip time: Smoothed 160 ms, Variance: 38
    Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
    Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
    Congestion window: Current: 1 KB

  FCP Session details <-----FCP session details
  Target FCID: 0x6c01e8 (S_ID of this session: 0x6c0202)
  pWWN: 21:00:00:20:37:62:c0:0c, nWWN: 20:00:00:20:37:62:c0:0c
  Session state: CLEANUP
  1 iSCSI sessions share this FC session
  Target: VT1
  Negotiated parameters
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
RcvDataFieldSize 1392 our_RcvDataFieldSize 1392
MaxBurstSize 0, EMPD: FALSE
Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
PDU: Command: 0, Response: 0
```

Use the **show fcns database** (and optionally **detail**) to display the Fibre Channel name server entry for the Fibre Channel N port created for iSCSI initiators in the SAN. See Example 4-12 and Example 4-13.

Example 4-12 Displays the FCNS Database Contents

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x020101      N     22:04:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w <---iSCSI
0x020102      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w initiator
0x0205d4      NL    21:00:00:04:cf:da:fe:c6 (Seagate)         scsi-fcp:target
0x0205d5      NL    21:00:00:04:cf:e6:e4:4b (Seagate)         scsi-fcp:target
...
Total number of entries = 10

VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xef0001      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w
Total number of entries = 1

VSAN 3:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xed0001      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w
Total number of entries = 1
```

Example 4-13 Displays the FCNS Database in Detail

```
switch# show fcns database detail
-----
VSAN:1      FCID:0x020101
-----
port-wwn (vendor)      :22:04:00:05:30:00:35:e1 (Cisco)
node-wwn                :22:03:00:05:30:00:35:e1
class                   :2,3
node-ip-addr            :10.2.2.12                <--- iSCSI initiator's IPv4 address
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1991-05.com.microsoft:oasis2-dell <--- iSCSI initiator's IQN
port-type               :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :22:01:00:05:30:00:35:de
hard-addr               :0x000000
-----
VSAN:1      FCID:0x020102
-----
port-wwn (vendor)      :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn                :22:01:00:05:30:00:35:e1
class                   :2,3
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

node-ip-addr      :10.2.2.11
ipa              :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type        :N
port-ip-addr     :0.0.0.0
fabric-port-wwn  :22:01:00:05:30:00:35:de
hard-addr        :0x000000
...
Total number of entries = 10
=====
-----
VSAN:2          FCID:0xef0001
-----
port-wwn (vendor) :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn          :22:01:00:05:30:00:35:e1
class             :2,3
node-ip-addr      :10.2.2.11
ipa              :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type        :N
port-ip-addr     :0.0.0.0
fabric-port-wwn  :22:01:00:05:30:00:35:de
hard-addr        :0x000000
Total number of entries = 1
...

```

Use the **show iscsi initiator configured** to display information about all the configured iSCSI initiators. Specifying the name shows information about the desired initiator. See Example 4-14.

Example 4-14 Displays Information About Configured Initiators

```

switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Member of vsans: 1, 2, 10
  Node WWN is 22:01:00:05:30:00:10:e1
  No. of PWWN: 5
    Port WWN is 22:04:00:05:30:00:10:e1
    Port WWN is 22:05:00:05:30:00:10:e1
    Port WWN is 22:06:00:05:30:00:10:e1
    Port WWN is 22:07:00:05:30:00:10:e1
    Port WWN is 22:08:00:05:30:00:10:e1

iSCSI Node name is 10.10.100.199
  Member of vsans: 1, 5
  Node WWN is 22:03:00:05:30:00:10:e1
  No. of PWWN: 4
    Port WWN is 22:00:00:05:30:00:10:e1
    Port WWN is 22:09:00:05:30:00:10:e1
    Port WWN is 22:0a:00:05:30:00:10:e1
    Port WWN is 22:0b:00:05:30:00:10:e1

User Name for Mutual CHAP: testuser

```


Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the Fibre Channel targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See Example 4-15.

Example 4-15 Displays Exported Targets

```
switch# show iscsi virtual-target
target: VT1
  * Port WWN 21:00:00:20:37:62:c0:0c
  Configured node
  all initiator permit is enabled

target: VT2
  Port WWN 21:00:00:04:cf:4c:52:c1
  Configured node
  all initiator permit is disabled
target: iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
  Port WWN 21:00:00:20:37:a6:be:32 , VSAN 1
  Auto-created node
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See Example 4-16.

Example 4-16 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdfg

username:user2
secret:cshadhdsadadjajdjas
```

Displaying iSLB VRRP Information

Use the **show islb vrrp summary vr** command to display VRRP load balancing information.

```
switch# show islb vrrp summary vr 30
```

```

-- Groups For Load Balance --
-----
VR Id          VRRP Address Type          Configured Status
-----
30             IPv4                        Enabled

-- Interfaces For Load Balance --
-----
VR Id          VRRP IP          Switch WWN          Ifindex          Load
-----
30  192.168.30.40  20:00:00:0d:ec:02:cb:00  GigabitEthernet3/1  2000
30  192.168.30.40  20:00:00:0d:ec:02:cb:00  GigabitEthernet3/2  2000
30  192.168.30.40  20:00:00:0d:ec:0c:6b:c0  GigabitEthernet4/1  2000
M 30  192.168.30.40  20:00:00:0d:ec:0c:6b:c0  GigabitEthernet4/2  1000
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying Pending iSLB Configuration Changes

You can display the pending configuration changes using the **show islb pending** command.

```
switch# show islb pending
iscsi initiator idle-timeout 10
islb initiator ip-address 10.1.1.1
static pWWN 23:01:00:0c:85:90:3e:82
static pWWN 23:06:00:0c:85:90:3e:82
username test1
islb initiator ip-address 10.1.1.2
static nWWN 23:02:00:0c:85:90:3e:82
```

You can display the differences between the pending configuration and the current configuration using the **show islb pending-diff** command.

```
switch# show islb pending-diff
+iscsi initiator idle-timeout 10
islb initiator ip-address 10.1.1.1
+ static pWWN 23:06:00:0c:85:90:3e:82
+islb initiator ip-address 10.1.1.2
+ static nWWN 23:02:00:0c:85:90:3e:82
```

Displaying iSLB CFS Status

You can display the iSLB CFS status using the **show islb session status** command.

```
switch# show islb status
iSLB Distribute is enabled
iSLB CFS Session exists
```

Displaying iSLB CFS Distribution Session Status

You can display the status of the iSLB CFS distribution session using the **show islb cfs-session status** command.

```
switch# show islb cfs-session status
last action          : fabric distribute enable
last action result   : success
last action failure cause : success
```

Displaying iSLB CFS Merge Status

You can display the iSLB CFS merge status using the **show islb merge status** command.

```
switch# show islb merge status
Merge Status: Success
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Verifying iSNS Client Configuration

Use the **show isns profile** command to view configured iSNS profiles. Profile ABC has two portals registered with the iSNS server. Each portal corresponds to a particular interface. Profile XYZ has a specified iSNS server, but does not have any tagged interfaces configured (see Example 4-17 and Example 4-18).

Example 4-17 Displays Information for Configured iSNS Profiles

```
switch# show isns profile
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204

iSNS profile name XYZ
iSNS Server 10.10.100.211
```

Example 4-18 Displays a Specified iSNS Profile

```
switch# show isns profile ABC
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204
```

Use the **show isns profile counters** command to view all configured profiles with the iSNS PDU statistics for each tagged interface (see Example 4-19 and Example 4-20).

Example 4-19 Displays Configured Profiles with iSNS Statistics

```
switch# show isns profile counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
  Output 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204

iSNS profile name XYZ
tagged interface port-channel 2
iSNS statistics
  Input 30 pdus (registration/deregistration pdus only)
    Reg pdus 29, Dereg pdus 1
  Output 30 pdus (registration/deregistration pdus only)
    Reg pdus 29, Dereg pdus 1
iSNS Server 10.1.4.218
```

Example 4-20 Displays iSNS Statistics for a Specified Profile

```
switch# show isns profile ABC counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

Reg pdus 37, Dereg pdus 17
Output 54 pdus (registration/deregistration pdus only)
Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204

```

Use the **show isns** command to view all objects registered on the iSNS server and specified in the given profile (see Example 4-21).

Example 4-21 Displays iSNS Queries

```

switch# show isns query ABC gigabitethernet 2/3
iSNS server: 10.10.100.204
Init: iqn.1991-05.com.w2k
  Alias: <MS SW iSCSI Initiator>
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03.210000203762fa34
  nWWN: 200000203762fa34

```

Use the **show interface** command to view the iSNS profile to which an interface is tagged (see Example 4-22).

Example 4-22 Displays Tagged iSNS Interfaces

```

switch# show interface gigabitethernet 2/3
GigabitEthernet2/3 is up
Hardware is GigabitEthernet, address is 0005.3000.ae94
Internet address is 10.10.100.201/24
MTU 1500 bytes
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
iSNS profile ABC
^^^^^^^^^^^^^^^^
5 minutes input rate 112 bits/sec, 14 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1935 packets input, 132567 bytes
  4 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
1 packets output, 42 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors

```

Verifying the iSNS Server Configuration

Use the **show isns config** command to view the ESI interval and the summary information about the iSNS database contents (see Example 4-23).

Example 4-23 Displays the iSNS Server Configuration of ESI Interval and Database Contents

```

switch# show isns config
Server Name: switch1(Cisco Systems) Up since: Fri Jul 30 04:08:16 2004
  Index: 1   Version: 1   TCP Port: 3205
  fabric distribute (remote sync): ON
  ESI
    Non Response Threshold: 5 Interval(seconds): 60
  Database contents
    Number of Entities: 2

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Number of Portals: 3
Number of iSCSI devices: 4
Number of Portal Groups: 0
```

Use the **show isns database** command to view detailed information about the contents of the iSNS database (see Example 4-24 through Example 4-27). This command displays the full iSNS database giving all the entities, nodes, and portals registered in the database. This command without options only displays explicitly registered objects. The asterisk next to the VSAN ID indicates that the iSCSI node is in the default zone for that VSAN.

Example 4-24 Displays Explicitly Registered Objects

```
switch# show isns database
Entity Id: dp-204
      Index: 2                Last accessed: Fri Jul 30 04:08:46 2004

iSCSI Node Name: iqn.1991-05.comdp-2041
      Entity Index: 2
      Node Type: Initiator(2)      Node Index: 0x1
      SCN Bitmap: OBJ_UPDATED|OBJ ADDED|OBJ REMOVED|TARGET&SELF
      Node Alias: <MS SW iSCSI Initiator>

      VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2      TCP Port: 4179
      Entity Index: 2      Portal Index: 1
      ESI Interval: 0      ESI Port: 4180      SCN Port: 4180
```

Example 4-25 displays information about both virtual and registered iSCSI initiators/targets.

Example 4-25 Displays the Full Database with Both Registered and Configured Nodes and Portals

```
switch# show isns database full
Entity Id: isns.entity.mds9000
      Index: 1                Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
      Entity Index: 1
      Node Type: Target(1)      Node Index: 0x80000001
      WWN(s):
          22:00:00:20:37:39:dc:45
      VSANS:

iSCSI Node Name: iqn.isns-first-virtual-target
      Entity Index: 1
      Node Type: Target(1)      Node Index: 0x80000002

      VSANS:

iSCSI Node Name: iqn.com.cisco.disk2
      Entity Index: 1
      Node Type: Target(1)      Node Index: 0x80000003
      WWN(s):
          22:00:00:20:37:39:dc:45

      VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
      Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
      Entity Index: 1      Portal Index: 5

Entity Id: dp-204
      Index: 2                Last accessed: Fri Jul 30 04:08:46 2004
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
iSCSI Node Name: iqn.1991-05.com.microsoft:dp-2041
  Entity Index: 2
  Node Type: Initiator(2)      Node Index: 0x1
  SCN Bitmap: OBJ_UPDATED|OBJ_ADDED|OBJ_REMOVED|TARGET&SELF
  Node Alias: <MS SW iSCSI Initiator>

  VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2      TCP Port: 4179
  Entity Index: 2      Portal Index: 1
  ESI Interval: 0      ESI Port: 4180      SCN Port: 4180
```

Example 4-26 displays the virtual targets entries on the current switch.



Note

The **local** option is only available for virtual targets.

Example 4-26 Displays the Virtual Target Information in the Local Switch

```
switch# show isns database virtual-targets local
Entity Id: isns.entity.mds9000
  Index: 1      Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000001
  WWN(s):
    22:00:00:20:37:39:dc:45

  VSANS:
iSCSI Node Name: iqn.isns-first-virtual-target
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000002

  VSANS:
iSCSI Node Name: iqn.com.cisco.disk2
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000003
  WWN(s):
    22:00:00:20:37:39:dc:45

  VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
  Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
  Entity Index: 1      Portal Index: 5
```

Example 4-27 provides the virtual target information for a specific remote switch. The remote switch is specified using the switch ID (the WWN of the switch).

Example 4-27 Displays Virtual Target for a Specified Switch

```
switch# show isns database virtual-targets switch 20:00:00:0d:ec:01:04:40
Entity Id: isns.entity.mds9000
  Index: 1      Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: iqn.com.cisco.disk1
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000001
  WWN(s):
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

22:00:00:20:37:39:dc:45

VSANS:
iSCSI Node Name: iqn.isns-first-virtual-target
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000002

VSANS:
iSCSI Node Name: iqn.com.cisco.disk2
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000003
WWN(s):
    22:00:00:20:37:39:dc:45

VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
Entity Index: 1      Portal Index: 5

```

Use the **show isns node** command to display attributes of nodes registered with the iSNS server (see Example 4-28 through Example 4-30). If you do not specify any options, the server displays the name and node type attribute in a compact format; one per line.

Example 4-28 Displays Explicitly Registered Objects

```

switch# show isns node all
-----
iSCSI Node Name                                     Type
-----
iqn.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8      Target
...
iqn.com.cisco.disk1                                         Target
iqn.com.cisco.ipdisk                                        Target
iqn.isns-first-virtual-target                              Target
iqn.1991-05.cw22                                           Target
iqn.1991-05.cw53                                           Target

```

Example 4-29 Displays the Specified Node

```

switch# show isns node name iqn.com.cisco.disk1
iSCSI Node Name: iqn.com.cisco.disk1
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000001
WWN(s):
    22:00:00:20:37:39:dc:45
VSANS: 1

```

Example 4-30 Displays the Attribute Details for All Nodes

```

switch# show isns node all detail
iSCSI Node Name: iqn.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8f
Entity Index: 1
Node Type: Target(1)      Node Index: 0x30000003
Configured Switch WWN: 20:00:00:0d:ec:01:04:40
WWN(s):
    22:00:00:20:37:5a:6c:8f
VSANS: 1

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

...
iSCSI Node Name: iqn.com.cisco.disk1
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000001
  Configured Switch WWN: 20:00:00:0d:ec:01:04:40
  WWN(s):
    22:00:00:20:37:39:dc:45
  VSANS: 1

iSCSI Node Name: iqn.com.cisco.ipdisk
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000002
  Configured Switch WWN: 20:00:00:0d:ec:01:04:40
  WWN(s):
    22:00:00:20:37:5a:70:1a
  VSANS: 1

iSCSI Node Name: iqn.isns-first-virtual-target
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000003
  Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: iqn.parna.121212
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000004
  Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: iqn.parna.121213
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000005
  Configured Switch WWN: 20:00:00:0d:ec:01:04:40

```

Use the **show isns portal** command to display the attributes of a portal along with its accessible nodes (see Example 4-31 through Example 4-35). You can specify portals by using the switch WWN-interface combination or the IP address-port number combination.

Example 4-31 Displays the Attribute Information for All Portals

```

switch# show isns portal all
-----
IPAddress      TCP Port      Index          SCN Port      ESI  port
-----
192.168.100.5  3205          3              -             -
192.168.100.6  3205          5              -             -

```

Example 4-32 Displays Detailed Attribute Information for All Portals

```

switch# show isns portal all detail
Portal IP Address: 192.168.100.5      TCP Port: 3205
  Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
  Entity Index: 1      Portal Index: 5

```

Example 4-33 Displays Virtual Portals

```

switch# show isns portal virtual
-----
IPAddress      TCP Port      Index          SCN Port      ESI  port
-----

```


Send documentation comments to dcnm-san-docfeedback@cisco.com

```
-----
192.168.100.5    3205    3    -    -
192.168.100.6    3205    5    -    -
```

Example 4-34 Displays Virtual Portals for the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40
-----
IPAddress      TCP Port    Index      SCN Port    ESI  port
-----
192.168.100.5    3205      3          -          -
192.168.100.6    3205      5          -          -
```

Example 4-35 Displays Detailed Information for the Virtual Portals in the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40 detail
Portal IP Address: 192.168.100.5      TCP Port: 3205
  Entity Index: 1    Portal Index: 3
  Switch WWN: 20:00:00:0d:ec:01:04:40
  Interface: GigabitEthernet2/3

Portal IP Address: 192.168.100.6      TCP Port: 3205
  Entity Index: 1    Portal Index: 5
  Switch WWN: 20:00:00:0d:ec:01:04:40
  Interface: GigabitEthernet2/5
```

Use the **show isns entity** command to display the attributes of an entity along with the list of portals and nodes in that entity (see Example 4-36 through Example 4-40). If you do not specify any option, this command displays the entity ID and number of nodes or portals associated with the entity in a compact format; one per line.

Example 4-36 Displays All Registered Entries

```
switch1# show isns entity
-----
Entity ID                                           Last Accessed
-----
dp-204                                             Tue Sep  7 23:15:42 2004
```

Example 4-37 Displays All Entities in the Database

```
switch# show isns entity all
-----
Entity ID                                           Last Accessed
-----
isns.entity.mds9000                                Tue Sep  7 21:33:23 2004
dp-204                                             Tue Sep  7 23:15:42 2004
```

Example 4-38 Displays the Entity with the Specified ID

```
switch1# show isns entity id dp-204
Entity Id: dp-204
  Index: 2          Last accessed: Tue Sep  7 23:15:42 2004
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 4-39 Displays Detailed Information for All Entities in the Database

```
switch1# show isns entity all detail
Entity Id: isns.entity.mds9000
  Index: 1          Last accessed: Tue Sep  7 21:33:23 2004

Entity Id: dp-204
  Index: 2          Last accessed: Tue Sep  7 23:16:34 2004
```

Example 4-40 Displays Virtual Entities

```
switch# show isns entity virtual
Entity Id: isns.entity.mds9000
  Index: 1          Last accessed: Thu Aug  5 00:58:50 2004

Entity Id: dp-204
  Index: 2          Last accessed: Thu Aug  5 01:00:23 2004
```

Use the **show iscsi global config** command to display information about import targets (see Example 4-41 and Example 4-42).

Example 4-41 Displays the Import Target Settings for the Specified Switch

```
switch# show isns iscsi global config switch 20:00:00:05:ec:01:04:00
iSCSI Global configuration:
  Switch: 20:00:00:05:ec:01:04:00 iSCSI Auto Import: Enabled
```

Example 4-42 Displays the Import Target Settings for All Switches

```
switch# show isns iscsi global config all
iSCSI Global configuration:
  Switch: 20:00:44:0d:ec:01:02:40 iSCSI Auto Import: Enabled
```

Use the **show cfs peers** command to display CFS peers switch information about the iSNS application (see Example 4-43).

Example 4-43 Displays the CFS Peer Switch Information for the iSNS Application

```
switch# show cfs peers name isns

Scope      : Physical
-----
Switch WWN          IP Address
-----
20:00:00:00:ec:01:00:40  10.10.100.11  [Local]

Total number of entries = 1
```

Verifying Automatic iSNS Cloud Discovery Configuration

To verify the automatic iSNS cloud discovery configuration, use the **show cloud discovery config** command.

```
switch# show cloud discovery config
Auto discovery: Enabled
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Verifying Cloud Discovery Status

Use the `show cloud discovery status` command to verify the status of the cloud discovery operation.

```
switch# show cloud discovery status
Discovery status: Succeeded
```

Verifying Cloud Discovery Membership

Use the `show cloud membership all` command to verify the cloud membership for the switch.

```
switch# show cloud membership all
Cloud 2
  GigabitEthernet1/5[20:00:00:0d:ec:02:c6:c0] IP Addr 10.10.10.5
  GigabitEthernet1/6[20:00:00:0d:ec:02:c6:c0] IP Addr 10.10.10.6
#members=2
```

Use the `show cloud membership unresolved` command to verify the unresolved membership on the switch.

```
switch# show cloud membership unresolved
Undiscovered Cloud
  No members
```

Displaying Cloud Discovery Statistics

Use the `show cloud discovery statistics` command to display the statistics for the cloud discovery operation.

```
switch# show cloud discovery statistics
Global statistics
  Number of Auto Discovery                = 1
  Number of Manual Discovery              = 0
  Number of cloud discovery (ping) messages sent = 1
  Number of cloud discovery (ping) success   = 1
```

Configuration Examples for iSCSI


This section provides three examples of iSCSI virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as an iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see Figure 4-17).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-17 Assigning iSCSI Node Names




```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
pWWN 28:00:01:02:03:04:05:06
```

112189

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 4-18).

Figure 4-18 Mapping LUNs to an iSCSI Node Name



```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```


112190

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 4-19).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-19 Mapping LUNs to Multiple iSCSI Node Names



112191

```

iscsi virtual-target name iqn.1987-02.com.cisco.target-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
  pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
  pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
  
```


Example of VSAN Membership for iSCSI Devices

Figure 4-20 provides an example of VSAN membership for iSCSI devices:

- iSCSI interface 1/1 is a member of VSAN Y.
- iSCSI initiator host A has explicit VSAN membership to VSAN X.
- Three iSCSI initiators (host A, host B, and host C) C connect to iSCSI interface 1/1.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-20 VSAN Membership for iSCSI Interfaces



Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so they will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

Example 4-44 and Example 4-45 are based on the following configurations:

- GigabitEthernet2/1.441 is the VRRP master interface for Switch1.
- GigabitEthernet2/2.441 is the VRRP backup interface for Switch1.
- GigabitEthernet1/1.441 is the VRRP backup interface for Switch2.
- GigabitEthernet1/2.441 is the VRRP backup interface for Switch2.

Example 4-44 Load Distribution with the Default Metric

The follow example output shows the initial load distribution for three initiators with the default load metric value:

```
switch# show islb vrrp summary
```

```
...
```

```
-----
VR Id   VRRP IP           Switch WWN           Ifindex              Load
-----
M 1     10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 0
  1     10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 1000
  1     10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 1000
  1     10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 1000
-- Initiator To Interface Assignment --
```

```
-----
Initiator                VR Id VRRP IP           Switch WWN           Ifindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
```

The following example output shows load distribution for four initiators. The interface load metric value for the master interface changed from 0 to 1000.

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch# show islb vrrp summary
...
-----
VVR Id   VRRP IP           Switch WWN           Ifindex             Load
-----
M 1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 1000
  1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 1000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 1000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 1000
      -- Initiator To Interface Assignment --
-----
Initiator           VR Id VRRP IP           Switch WWN           Ifindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

The following example output shows load distribution for nine initiators. The interface load metric values for the backup interfaces have changed.

```
switch# show islb vrrp summary
...
-----
VVR Id   VRRP IP           Switch WWN           Ifindex             Load
-----
M 1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 1000
  1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 2000
      -- Initiator To Interface Assignment --
-----
Initiator           VR Id VRRP IP           Switch WWN           Ifindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
iqn.cisco.test-linux.init4 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init5 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init6 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init7 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init8 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
```

Example 4-45 Load Distribution with the Metric Set to 3000 on One Initiator

The following example output shows the initial load distribution for three initiators with one initiator having load metric of 3000 and the remaining initiator with the default metric value:

```
switch# show islb vrrp summary
...
-----
VVR Id   VRRP IP           Switch WWN           Ifindex             Load
-----
M 1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 0
  1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 1000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 1000
      -- Initiator To Interface Assignment --
-----
Initiator           VR Id VRRP IP           Switch WWN           Ifindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
```

The follow example output shows load distribution for four initiators. The interface load metric value for the master interface changed from 0 to 1000.

```
switch# show islb vrrp summary
```

```
...
-----
VVR Id   VRRP IP           Switch WWN           Iindex              Load
-----
M 1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 1000
  1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 1000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 1000
      -- Initiator To Interface Assignment --
-----
Initiator                VR Id VRRP IP           Switch WWN           Iindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init3 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

The following example output shows load distribution for nine initiators. The interface load metric values for the backup interfaces have changed.

```
switch# show islb vrrp summary
```


```
...
-----
VVR Id   VRRP IP           Switch WWN           Iindex              Load
-----
M 1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441 2000
  1      10.10.122.115    20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441 3000
  1      10.10.122.115    20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441 3000
      -- Initiator To Interface Assignment --
-----
Initiator                VR Id VRRP IP           Switch WWN           Iindex
-----
iqn.cisco.test-linux.init0 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/2.441
iqn.cisco.test-linux.init1 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init2 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init3 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
iqn.cisco.test-linux.init4 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init5 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init6 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/1.441
iqn.cisco.test-linux.init7 1 10.10.122.115 20:00:00:0c:ce:5c:5b:c0 GigabitEthernet1/2.441
iqn.cisco.test-linux.init8 1 10.10.122.115 20:00:00:0b:5f:3c:01:80 GigabitEthernet2/1.441
```

Example of an iSNS Server

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. Figure 4-21 provides an example of this scenario.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-21 Using iSNS Servers in the Cisco MDS Environment



In Figure 4-21, iqn.host1 and iqn.host2 are iSCSI initiators. P1 and P2 are Fibre Channel targets. The two initiators are in different zones: Zone 1 consists of iqn.host1 and target P1, and Zone 2 consists of iqn.host2 and target P2. iSNS server functionality is enabled on both switches, SW-1 and SW-2. The registration process proceeds as follows:

1. Initiator iqn.host1 registers with SW-1, port GigabitEthernet2/1.
2. Initiator iqn.host2 registers with SW-2, port GigabitEthernet3/1.
3. Initiator iqn.host1 issues an iSNS query to SW-1 to determine all accessible targets.
4. The iSNS server in turn queries the Fibre Channel name server (FCNS) to obtain a list of devices that are accessible (that is, in the same zone) by the query originator. This query yields only P1.
5. The iSNS server then queries its own database to convert the Fibre Channel devices to the corresponding iSCSI targets. This is based on the iSCSI configuration, such as virtual-target and its access control setting or whether the dynamic Fibre Channel target import feature is enabled or disabled.
6. The iSNS server sends a response back to the query initiator. This response contains a list all iSCSI portals known to the iSNS server. This means iqn.host1 can choose to log in to target P1 through either SW-1 (at GigabitEthernet 2/1) or SW-2 (at GigabitEthernet 3/1).
7. If the initiator chooses to log in to SW-1 and later that port becomes inaccessible (for example, GigabitEthernet 2/1 goes down), the initiator has the choice to move to connect to target P1 through port GigabitEthernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

iSCSI Transparent Mode Initiator Example

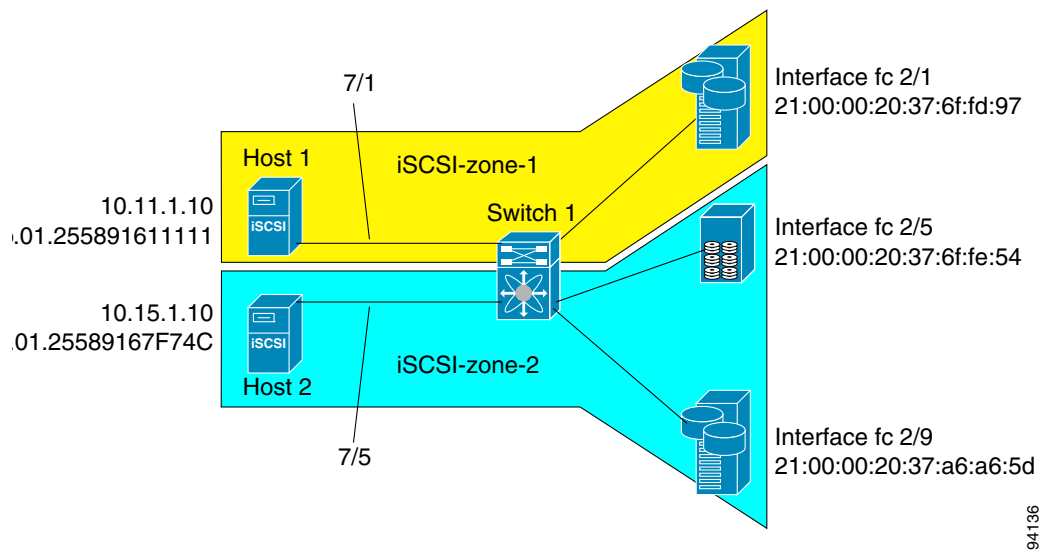
This examples assumes the following configuration (see Figure 4-22):

- No LUN mapping or LUN masking or any other access control for hosts on the target device
- No iSCSI login authentication (that is, login authentication set to none)
- The topology is as follows:
 - iSCSI interface 7/1 is configured to identify initiators by IP address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- iSCSI interface 7/5 is configured to identify initiators by node name.
- The iSCSI initiator host 1 with IPv4 address 10.11.1.10 and name iqn.1987-05.com.cisco:01.255891611111 connects to IPS port 7/1 is identified using IPv4 address (host 1 = 10.11.1.10).
- The iSCSI initiator host 2 with IPv4 address 10.15.1.10 and node name iqn.1987-05.com.cisco:01.25589167f74c connects to IPS port 7/5.


Figure 4-22 iSCSI Transparent Mode Initiator



94196

Detailed Steps

To configure scenario 1 (see Figure 4-22), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts in Cisco MDS switches.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.
- ```
switch(config)# interface gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shutdown
```
-  **Note** Host 2 is connected to this port.
- 
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address, and enable the interface.
- ```
switch(config)# interface iscsi 7/1
switch(config-if)# switchport initiator id ip-address
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch(config-if)# no shut
```

- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.

```
switch(config)# interface gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shutdown
```

- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.

```
switch(config)# interface iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shutdown
```



Note Host 1 is connected to this port.

- Step 7** Verify the available Fibre Channel targets (see Figure 4-22).

```
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 ip-address 10.11.1.10
```

- Step 9** Create a zone named *iscsi-zone-2* with host 2 and two Fibre Channel targets in it.



Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

```
switch(config)# zone name iscsi-zone-2 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fe:54
switch(config-zone)# member pwn 21:00:00:20:37:a6:a6:5d
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c
```

- Step 10** Create a zone set and add the two zones as members.

```
switch(config)# zoneset name zoneset-iscsi vsan 1
switch(config-zoneset)# member iscsi-zone-1
switch(config-zoneset)# member iscsi-zone-2
```

- Step 11** Activate the zone set.

```
switch(config)# zoneset activate name zoneset-iscsi vsan 1
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 12 Display the active zone set.



Note The iSCSI hosts are not connected so they do not have an FC ID yet.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
zone name iscsi-zone-1 vsan 1
* fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97] <-----Target
  symbolic-nodename 10.11.1.10 <-----iSCSI host (host 1, not online)

zone name iscsi-zone-2 vsan 1
* fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54] <-----Target
* fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d] <-----Target
  symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c <-iSCSI host (host 2, not online)
```

Step 13 Bring up the iSCSI hosts (host 1 and host 2).

Step 14 Show all the iSCSI sessions (use the **detail** option for detailed information).

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.25589167f74c <-----Host 2
Initiator ip addr (s): 10.15.1.11
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54
```



Note The last part of the auto-created target name is the Fibre Channel target's pWWN.

```
VSAN 1, ISID 00023d000001, Status active, no reservation

Session #2
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d
VSAN 1, ISID 00023d000001, Status active, no reservation

Initiator 10.11.1.10 <-----Host 1
Initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97
VSAN 1, ISID 00023d000001, Status active, no reservation
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.25589167f74c <-----
  Initiator ip addr (s): 10.15.1.11
  iSCSI alias name: oasis11.cisco.com
  Node WWN is 20:02:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:03:00:0b:fd:44:68:c2 (dynamic)
    Interface iSCSI 7/5, Portal group tag: 0x304
    VSAN ID 1, FCID 0x6d0300

iSCSI Node name is 10.11.1.10 <-----
  iSCSI Initiator name: iqn.1987 - 05.com.cisco:01.e41695d16b1a
  iSCSI alias name: oasis10.cisco.com
  Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:05:00:0b:fd:44:68:c2 (dynamic)
    Interface iSCSI 7/1, Portal group tag: 0x300
    VSAN ID 1, FCID 0x6d0301
```

Host 2: Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Host 1: Initiator ID based on IPv4 address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FC IDs are resolved.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x6d0301 [symbolic-nodename 10.11.1.10] <-----

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d]
    * fcid 0x6d0300 [symbolic-nodename
iqn.1987-05.com.cisco:01.25589167f74c] <-----
```

FC ID resolved for host 1

FC ID for host 2

Step 17 The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE
-----
0x6d0001      NL    21:00:00:20:37:6f:fd:97 (Seagate)         scsi-fcp:target
0x6d0101      NL    21:00:00:20:37:6f:fe:54 (Seagate)         scsi-fcp:target
0x6d0201      NL    21:00:00:20:37:a6:a6:5d (Seagate)         scsi-fcp:target
0x6d0300      N     20:03:00:0b:fd:44:68:c2 (Cisco)             scsi-fcp:init isc..w
0x6d0301      N     20:05:00:0b:fd:44:68:c2 (Cisco)             scsi-fcp:init isc..w
```

Send documentation comments to dcnm-san-docfeedback@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     :

symbolic-node-name
:ign.1987-05.com.cisco:01.25589167f74c<-----
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn        :21:91:00:0b:fd:44:68:c0
hard-addr               :0x000000
Total number of entries = 1
```

IPv4 address of the
iSCSI host
iSCSI gateway node

iSCSI initiator ID is
based on the registered
node name

```
switch# show fcns database fcid 0x6d0301 detail vsan 1
-----
VSAN:1      FCID:0x6d0301
-----
port-wwn (vendor)      :20:05:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:04:00:0b:fd:44:68:c2
class                   :2,3
node-ip-addr            :10.11.1.10
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name     :

symbolic-node-name     :10.11.1.10  <-----
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn        :21:81:00:0b:fd:44:68:c0
hard-addr               :0x000000
```

iSCSI gateway node

iSCSI initiator ID is
based on the IPv4
address registered in
symbolic-node-name
field

To configure this example (see Figure 4-22), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts in Cisco MDS switches.
- In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - Select **none** from the AuthMethod drop-down menu in the Information pane.
 - Click the **Apply Changes** icon.
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
- In Device Manager, click **IP > iSCSI**.
 - Click the **Targets** tab.
 - Check the **Dynamically Import FC Targets** check box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- d. Click **Apply**.
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
 - Select the **IP Address** tab in the Information pane and click **Create Row**.
 - Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.
 - Click **Create**.
 - Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.
 - Click the **Apply Changes** icon.



Note Host 2 is connected to this port.

- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address, and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
 - Click the **iSCSI** tab in the Information pane.
 - Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
 - In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
 - Click the **iSCSI** tab.
 - Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.
 - Click **Apply**.
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with an IPv4 address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
 - Click the **IP Address** tab in the Information pane and click **Create Row**.
 - Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
 - Click **Create**.
 - Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
 - Click the **Apply Changes** icon.
- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
 - Click the **iSCSI** tab in the Information pane.
 - Select **name** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
 - In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
 - Click the **iSCSI** tab.
 - Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
 - Click **Apply**.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets.

- a. In Device Manager, Choose **FC > Name Server**.
- b. Click the **General** tab.

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.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zones** folder in the left navigation pane and click **Insert**.
- d. Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- e. Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.
- f. Set the ZoneBy radio button to **WWN**.
- g. Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97) and click **Add**.
- h. Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- i. Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.

Step 9 Create a zone named iscsi-zone-2 with host 2 and two Fibre Channel targets in it.



Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** from the main menu.
- b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zones** folder in the left navigation pane and click **Insert**.
- d. Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- e. Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.
- f. Set the ZoneBy radio button to **WWN**.
- g. Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5). and click **Add**.
- h. Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d). and click **Add**.
- i. Set the ZoneBy radio button to **iSCSI name**.
- j. Set the Port Name field to the symbolic name for host 2 (iqn.1987-05.com.cisco:01.25589167f74c) and click **Add**.

Step 10 Create a zone set, add the two zones as members, and activate the zone set.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Note iSCSI interface is configured to identify all hosts based on node name.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
 - b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
 - c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.
 - d. Set the Zoneset Name to **zonset-iscsi** and click **OK**.
 - e. Click on the **zonset-iscsi** folder and click **Insert**.
 - f. Set the Zone Name field to **iscsi-zone-1** and click **OK**.
 - g. Set the Zone Name field to **iscsi-zone-2** and click **OK**.
 - h. Click **Activate** to activate the new zone set.
 - i. Click **Continue Activation** to finish the activation.
- Step 11** Bring up the iSCSI hosts (host 1 and host 2).
- Step 12** Show all the iSCSI sessions.
- a. In Device Manager, choose **Interfaces > Monitor > Ethernet**.
 - b. Click the **iSCSI connections** tab to show all the iSCSI sessions.
 - c. In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.
 - d. Click **Details**.
- Step 13** In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators
- Step 14** In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** to view the active zone set. The iSCSI initiators' FC IDs are resolved.
- Step 15** In Device Manager, Choose **FC > Name Server**. The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.
- Step 16** In Device Manager, Choose **FC > Name Server**.
- Step 17** Click the **Advanced** tab. Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.
-


Target Storage Device Requiring LUN Mapping Example

This example scenario 2 assumes the following configuration (see Figure 4-23):

- Access control is based on Fibre Channel zoning.
- There is target-based LUN mapping or LUN masking.
- There is no iSCSI authentication (none).
- The iSCSI initiator is assigned to different VSANs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-23 Target Storage Device with LUN Mapping



94137

Detailed Steps

To configure scenario 2 (see Figure 4-23), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.
- ```
switch(config)# interface gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shutdown
```
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address and enable the interface.
- ```
switch(config)# interface iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shutdown
```
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.
- ```
switch(config)# interface gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shutdown
```
- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.
- ```
switch(config)# interface iscsi 7/5
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shutdown
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 7 Add static configuration for each iSCSI initiator.

```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a <-----Host 2
switch(config-iscsi-init)# static pwwn system-assign 1
switch(config-iscsi-init)# static nwwn system-assign

switch(config)# iscsi initiator ip address 10.15.1.11 <-----Host 1
switch(config-iscsi-init)# static pwwn system-assigned 1
switch(config-iscsi-init)# vsan 2
```



Note Host 1 is configured in VSAN 2.

Step 8 View the configured WWNs.



Note The WWNs are assigned by the system. The initiators are members of different VSANs.

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
  Member of vsans: 1
  Node WWN is 20:03:00:0b:fd:44:68:c2
  No. of PWWN: 1
  Port WWN is 20:02:00:0b:fd:44:68:c2

iSCSI Node name is 10.15.1.11
  Member of vsans: 2
  No. of PWWN: 1
  Port WWN is 20:06:00:0b:fd:44:68:c2
```

Step 9 Create a zone with host 1.

```
switch(config)# zone name iscsi-zone-1 vsan 1
```

Step 10 Add three members to the zone named *iscsi-zone-1*.



Note Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN, can be used. In this case, the pWWN is persistent.

- The following command is based on the symbolic node name.

```
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
```

- The following command is based on the persistent pWWN assigned to the initiator. You can obtain the pWWN from the **show iscsi initiator** output.

```
switch(config-zone)# member pwwn 20:02:00:0b:fd:44:68:c2
```

Step 11 Create a zone with host 2 and two Fibre Channel targets.



Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

```
switch(config)# zone name iscsi-zone-2 vsan 2
```

Step 12 Activate the zone set in VSAN 2.

```
switch(config)# zoneset activate name iscsi-zoneset-v2 vsan 2
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Zoneset activation initiated. check zone status
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
    * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
      pwwn 20:06:00:0b:fd:44:68:c2 <-----Host is not online
```

Step 13 Start the iSCSI clients on both hosts and verify that sessions come up.

Step 14 Display the iSCSI sessions to verify the Fibre Channel target and the configured WWNs.

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.e41695d16b1a
  Initiator ip addr (s): 10.11.1.10
  Session #1
    Discovery session, ISID 00023d000001, Status active

  Session #2
    Target
    iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97<---- To Fibre Channel 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 iscw <--- iSCSI initiator in name server
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 17 Verify the details of the iSCSI initiator's FC ID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
-----
VSAN:1      FCID:0x680102
-----
port-wwn (vendor)      :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:03:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr           :10.11.1.10
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:81:00:0b:fd:44:68:c0
        iSCSI alias name: oasis10.cisco.com
```

Step 18 Check the Fibre Channel name server.

```
switch# show fcns database vsan 1
VSAN 1:
-----
FCID      TYPE  PWWN                                (VENDOR) FC4-TYPE:FEATURE
-----
0x680001  NL   21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target
0x680102  N    20:02:00:0b:fd:44:68:c2 (Cisco)   scsi-fcp:init isc..w <----- iSCSI
                                         initiator in
                                         name server
```

Step 19 Verify the details of the iSCSI initiator's FC ID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
-----
VSAN:1      FCID:0x680102
-----
port-wwn (vendor)      :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:03:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr           :10.11.1.10
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:81:00:0b:fd:44:68:c0
hard-addr              :0x000000
```

Step 20 Verify that zoning has resolved the FC ID for the iSCSI client.

```
switch# show zoneset active vsan 1
zoneset name iscsi-zoneset-v1 vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x680001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x680102 [pwwn 20:02:00:0b:fd:44:68:c2]
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 21 Verify that the second initiator is connected to the two Fibre Channel targets in VSAN 2.

```
switch# show iscsi session initiator 10.15.1.11
Initiator 10.15.1.11
  Initiator name iqn.1987-05.com.cisco:01.25589167f74c
  Session #1
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <-- Session to
    VSAN 2, ISID 00023d000001, Status active, no reservation                first target

  Session #2
    Target iqn.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: iqn.1987-05.com.cisco:01.25589167f74c
  iSCSI alias name: oasis11.cisco.com

  Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <----- 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
                                                                              entry in
                                                                              name server

switch# show fcns database fcid 0x750200 detail vsan 2
-----
VSAN:2      FCID:0x750200
-----
port-wwn (vendor)      :20:06:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:04:00:0b:fd:44:68:c2
class                 :2,3
node-ip-addr           :10.15.1.11
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :10.15.1.11
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:91:00:0b:fd:44:68:c0
hard-addr              :0x000000
Total number of entries = 1
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
    * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]


    * fcid 0x750200 [pwwn 20:06:00:0b:fd:44:68:c2] <-----
```

**FC ID
resolved for
iSCSI
initiator**

To configure this example (see Figure 4-23), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - Select **none** from the AuthMethod drop-down menu in the Information pane.
 - Click the **Apply Changes** icon.
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
- In Device Manager, click **IP > iSCSI**.
 - Click the **Targets** tab.
 - Check the **Dynamically Import FC Targets** check box.
 - Click **Apply**.
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
 - Select the **IP Address** tab in the Information pane and click **Create Row**.
 - Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.
 - Click **Create**.
 - Click the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.
 - Click the **Apply Changes** icon.
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by their IP address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
 - Select the **iSCSI** tab in the Information pane.
 - Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
 - In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
 - Click the **iSCSI** tab.
 - Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- g.** Click **Apply**.
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
 - Click the **IP Address** tab in the Information pane and click **Create Row**.
 - Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
 - Click **Create**.
 - Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
 - Click the **Apply Changes** icon.
- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.
- In Cisco DCNM-SAN, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
 - Click the **iSCSI** tab in the Information pane.
 - Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
 - In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
 - Click the **iSCSI** tab.
 - Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
 - Click **Apply**.
- Step 7** Configure for static pWWN and nWWN for host 1.
- In Device Manager, choose **IP > iSCSI**.
 - Click the **Initiators** tab.
 - Check the **Node Address Persistent** and **Node Address System-assigned** check boxes the Host 1 iSCSI initiator.
 - Click **Apply**.
- Step 8** Configure for static pWWN for Host 2.
- In Device Manager, Choose **IP > iSCSI**.
 - Click the **Initiators** tab.
 - Right-click on the Host 2 iSCSI initiator and click Edit pWWN.
 - Select **1** from the System-assigned Num field and click **Apply**.
- Step 9** View the configured WWNs.
-  **Note** The WWNs are assigned by the system. The initiators are members of different VSANs.
- In Cisco DCNM-SAN, choose **End Devices > iSCSI** in the Physical Attributes pane.
 - Click the **Initiators** tab.
- Step 10** Create a zone for Host 1 and the iSCSI target in VSAN 1.

Send documentation comments to dcnm-san-docfeedback@cisco.com



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.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zones** folder in the left navigation pane and click **Insert**.
- d. Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- e. Select the iscsi-zone-1 folder in the left navigation pane and click **Insert**.
- f. Set the ZoneBy radio button to **WWN**.
- g. Set the Port WWN to the pWWN for the Fibre Channel target (that is, 21:00:00:20:37:6f:fd:97), and click **Add**.
- h. Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- i. Set the IP Address/Mask field to the IP Address for Host 1 iSCSI initiator (10.11.1.10) and click **Add**.



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.

Step 11 Create a zone set in VSAN 1 and activate it.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b. Select VSAN 1 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.
- d. Set the Zoneset Name to **zonset-iscsi-1** and click **OK**.
- e. Click on the **zonset-iscsi-1** folder and click **Insert**.
- f. Set the Zone Name field to **iscsi-zone-1** and click **OK**.
- g. Click **Activate** to activate the new zone set.
- h. Click **Continue Activation** to finish the activation.

Step 12 Create a zone with host 2 and two Fibre Channel targets.



Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.



Note iSCSI interface is configured to identify all hosts based on node name.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zones** folder in the left navigation pane and click **Insert**.
- d. Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- e. Select the **iscsi-zone-2** folder in the left navigation pane and click **Insert**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- f. Set the ZoneBy radio button to **WWN**.
- g. Set the Port WWN to the pWWN for one of the Fibre Channel targets (for example, 21:00:00:20:37:6f:fe:5) and click **Add**.
- h. Set the Port WWN to the pWWN for another of the Fibre Channel targets (for example, 21:00:00:20:37:a6:a6:5d) and click **Add**.
- i. Set the ZoneBy radio button to **iSCSI IP Address/Subnet**.
- j. Set the IP Address/Mask field to the IP Address for Host 2 iSCSI initiator (10.15.1.11) and click **Add**.

Step 13 Create a zone set in VSAN 2 and activate it.

- a. In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database**.
- b. Select VSAN 2 from the VSAN drop-down menu in the Edit Local Full Zone Database dialog box.
- c. Select the **Zoneset** folder in the left navigation pane and click **Insert**.
- d. Set the Zoneset Name to **zonset-iscsi-2** and click **OK**.
- e. Click on the **zonset-iscsi-2** folder and click **Insert**.
- f. Set the Zone Name field to **iscsi-zone-2** and click **OK**.
- g. Click **Activate** to activate the new zone set.
- h. Click **Continue Activation** to finish the activation.

Step 14 Start the iSCSI clients on both hosts.

Step 15 Show all the iSCSI sessions.

- a. In Device Manager, choose **Interface > Monitor > Ethernet** and select the **iSCSI connections** tab to show all the iSCSI sessions.
- b. In Device Manager, choose **IP > iSCSI** and select the **Session Initiators** tab.
- c. Click **Details**.

Step 16 In Cisco DCNM- SAN, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators.

Step 17 In Cisco DCNM-SAN, choose **Zones > Edit Local Full Zone Database** to view the active zone set. The iSCSI initiators' FC IDs are resolved.

Step 18 In Device Manager, choose **FC > Name Server**. The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

Step 19 In Device Manager, Choose **FC > Name Server**.

Step 20 Click the **Advanced** tab. Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

Field Descriptions for iSCSI

The following are the field descriptions for iSCSI.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Ethernet Interfaces iSCSI

Field	Description
Description	An alias name for the interface as specified by a network manager.
Speed	Operational speed.
PhysAddress	The interface's WWN.
Admin	The desired state of the interface.
Oper	The current operational state of the interface.
LastChange	When the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this value contains a N/A value.
PortVSAN	The VSAN that the interface belongs to.
ForwardingMode	Use Store and Forward if the HBA has problems with passthrough.
Initiator ID Mode	How the initiator is identified on this interface, either by its iSCSI name (name) or by its IP address (ipaddress).
Enable	The initiator proxy mode for this interface. If true, then all the initiators coming on this interface would use the initiator configuration provided by this interface. The initiator configuration include port WWN and node WWN.
Assignment	How the initiator proxy mode FC addresses are assigned. If auto, then the FC addresses are automatically assigned. If it is manual, then they have to be manually configured.
Port WWN	The Port FC address used by the initiators on this interface when the initiator proxy mode is on.
Node WWN	The Node FC address used by the initiators on this interface when the initiator proxy mode is on.

Ethernet Interfaces iSCSI TCP

Field	Description
Local Port	Local interface TCP port.
SACK	Indicates if the Selective Acknowledgement (SACK) option is enabled or not.
KeepAlive	The TCP keepalive timeout for this iSCSI interface. If the value is 0, the keepalive timeout feature is disabled.
MinTimeout	The TCP minimum retransmit time.
Max	The TCP maximum retransmissions.
SendBufferSize	The TCP send buffer size.
MinBandwidth	The TCP minimum bandwidth.
MaxBandwidth	The TCP maximum bandwidth.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Estimated Round Trip	The estimated round trip delay of network pipe used for B-D product computation. The switch can use this to derive the TCP window to advertise.
QoS	The TCP QoS code point.
PMTU Enable	Indicates if the Path MTU discovery option is enabled or not.
PMTU Reset Timeout	The PMTU reset timeout.
Connections Normal	The number of normal iSCSI connections.
Connections Discovered	The number of discovery iSCSI connections.
CWM Enable	If true, congestion window monitoring is enabled. If false, it is disabled.
CWM Burst Size	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.
Port	The local TCP port of this interface.

Ethernet Interface Monitor iSCSI Connections

Field	Description
RxBytes	Total number of bytes received on an iSCSI session.
TxBytes	Total number of bytes transmitted on an iSCSI session.
IPSEC	A collection of objects for iSCSI connection statistics.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.
CID	The iSCSI Connection ID for this connection.
State	The current state of this connection, from an iSCSI negotiation point of view. <ul style="list-style-type: none"> login- The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. full- A valid iSCSI login response with the final bit set has been sent or received. logout- A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. Note that the size of reported in bytes even though the negotiation is in 512 k blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of command PDUs transferred on this session.
PDU Response	The count of response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

iSCSI Initiator PWWN

Field	Description
Port WWN	The FC address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> normal—session is a normal iSCSI session discovery—session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

iSNS Details iSCSI Nodes

Field	Description
Name	The iSCSI name of the initiator or target associated with the storage node.
Type	The Node Type bit-map defining the functions of this iSCSI node, where 31 is a Target, 30 is an Initiator, 29 is a Control, and all others are reserved.
Alias	The Alias name of the iSCSI node.
ScnBitmap	The State Change Notification (SCN) bitmap for a node.
WWN Token	An optional globally unique 64-bit integer value that can be used to represent the world wide node name of the iSCSI device in a Fibre Channel fabric.
AuthMethod	The iSCSI authentication method enabled for this iSCSI Node.

iSCSI User

Field	Description
iSCSI User	The name of the iSCSI user.
Password	The password of the iSCSI user.

Edit iSCSI Advertised Interfaces

Field	Description
Num	The number of the iSCSI target.
Interface	The interface over which the target is to be advertised.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Additional References

For additional information related to implementing FCIPs, see the following section:

- [Related Document, page 4-136](#)
- [Standards, page 4-136](#)
- [RFCs, page 4-136](#)
- [MIBs, page 4-136](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-ISCI-GW-MIB • CISCO-ISCSI-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



Configuring IP Services

This chapter includes the following topics:

- [Information About IP Services, page 5-1](#)
- [Guidelines and Limitations, page 5-7](#)
- [Default Settings, page 5-8](#)
- [Configuring IP Services, page 5-8](#)
- [Configuring Multiple VSANs, page 5-14](#)
- [Configuring VRRP, page 5-16](#)
- [Verifying IP Services Configuration, page 5-25](#)
- [Configuration Examples for IP Services, page 5-30](#)
- [Field Descriptions for IP Services, page 5-33](#)
- [Additional References, page 5-43](#)

Information About IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMSs):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on 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.

Text Part Number:

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

For information about configuring IPv6, see [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)


This section includes the following topics:

- [Traffic Management Services, page 5-2](#)
- [Management Interface Configuration, page 5-3](#)
- [About the Default Gateway, page 5-3](#)
- [IPv4 Default Network Configuration, page 5-4](#)
- [IPFC, page 5-5](#)
- [About IPv4 Static Routes, page 5-5](#)
- [About Overlay VSANs, page 5-5](#)
- [About VRRP, page 5-5](#)
- [DNS Server Configuration, page 5-7](#)

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running the IP protocol over an Fibre Channel 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 as shown in [Figure 5-1](#).

Figure 5-1 Management Access to Switches



Send documentation comments to dcnm-san-docfeedback@cisco.com

Management Interface Configuration

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure IP version 4 (IPv4) parameters (IP address, subnet mask) or an IP version 6 (IPv6) address and prefix length so that the switch is reachable. For information on configuring IPv6 addresses, see [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.



Note

The port on the Ethernet switch to which the MDS management interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should be disabled. This helps avoid the delay in the MDS management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in Cisco IOS or the **set port host** command in the Catalyst OS. Refer to the configuration guide for your Ethernet switch.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

About the Default Gateway

You can configure a default gateway IPv4 address on your Cisco MDS 9000 Family switch.

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes (IP default network, destination prefix, and destination mask, and next hop address). If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled.

The default gateway IPv4 address should be configured along with the IPv4 static routing attributes commands ([IP default network, destination prefix, and destination mask, and next hop address](#)).



Tip

If you configure the static route IP forwarding and the default-network details, these IPv4 addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

Use the **ip default-gateway** command to configure the IP address for a switch's default gateway and the **show ip route** command to verify that the IPv4 address for the default gateway is configured.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IPv4 Default Network Configuration

If you assign the IPv4 default network address, the switch considers routes to that network as the last resort. If the IPv4 default network address is not available, the switch uses the IPv4 default gateway address. For every network configured with the IPv4 default network address, the switch flags that route as a candidate default route, if the route is available.




Tip

If you configure the static route IP forwarding and the default network details, these IPv4 addresses will be used regardless of the default gateway being enabled or disabled. If these IPv4 addresses are configured and not available, the switch will fall back to using the default gateway IPv4 address, if you have configured it. Be sure to configure IPv4 addresses for all entries in the switch if you are using IPv4.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IPv4 address of the gateway switch as shown in Figure 5-2.

Figure 5-2 Overlay VSAN Functionality



In Figure 5-1, switch A has the IPv4 address 1.12.11.1, switch B has the IPv4 address 1.12.11.2, switch C has the IPv4 address 1.12.11.3, and switch D has the IPv4 address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IPv4 address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IPv4 address (1.12.11.1) in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet, 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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IPFC

IPFC provides IP forwarding on in-band switch management over a Fibre Channel interface (rather than out-of-band using the Gigabit Ethernet mgmt 0 interface). You can use IPFC to specify that IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.

Once the VSAN interface is created, you can specify the IP address for that VSAN. You can assign an IPv4 address or an IPv6 address.

**Note**

See the [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces”](#) for information about configuring IPv6 on the Cisco MDS 9000 Family switches.

About IPv4 Static Routes

Static routing is a mechanism to configure IPv4 routes on the switch. You can configure more than one static route.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IPv4 routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

If your network configuration does not need an external router, you can configure IPv4 static routing on your MDS switch.

**Note**

For information about IPv6 static routing, see the [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

About Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

About VRRP

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following features:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.

Send documentation comments to dcnm-san-docfeedback@cisco.com


- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and the draft-ietf-vrrp-ipv6 specification.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with different virtual router IP mapping.
- Both IPv4 and IPv6 is supported.
- The management interface (mgmt 0) supports only one virtual router group. All other interfaces each support up to seven virtual router groups, including both IPv4 and IPv6 combined. Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.



Note If you are using IPv6, you must either configure an IPv6 address on the interface or enable IPv6 on the interface. For more information about IPv6, see [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

In Figure 5-3, switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have an IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches do not have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

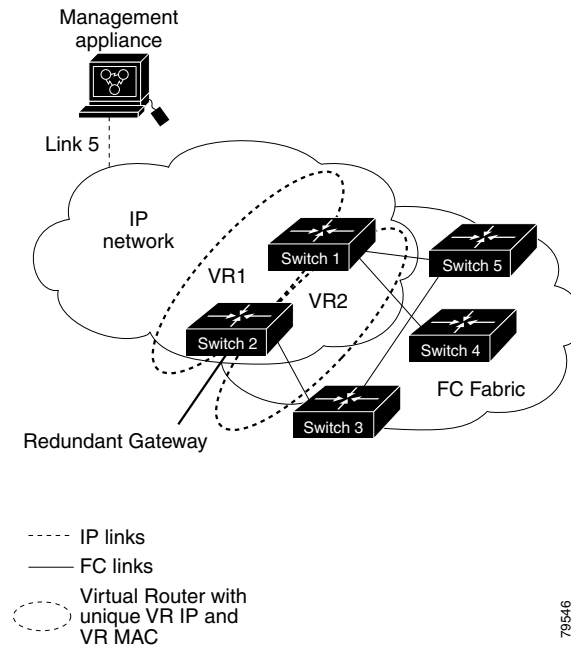
Figure 5-3 VRRP Functionality



In Figure 5-4, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 5-4 Redundant Gateway



DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts because of one of the following reasons:

- The IP address or the switch name is wrongly configured.
- The DNS server is not reachable because external reasons (reasons beyond our control).



Note

When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

Guidelines and Limitations

Follow these guidelines to configure IPFC:

1. Create the VSAN to use for in-band management, if necessary.
2. Configure an IPv4 address and subnet mask for the VSAN interface.
3. Enable IPv4 routing.
4. Verify connectivity.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Default Settings

Table 5-1 lists the default settings for DNS features.

Table 5-1 *Default DNS Settings*

Parameters	Default
Domain lookup	Disabled
Domain name	Disabled
Domains	None
Domain server	None
Maximum domain servers	6

Table 5-2 lists the default settings for VRRP features.

Table 5-2 *Default VRRP Settings*

Parameters	Default
Virtual router state	Disabled
Maximum groups per VSAN	255
Maximum groups per Gigabit Ethernet port	7
Priority preemption	Disabled
Virtual router priority	100 for switch with secondary IP addresses 255 for switches with the primary IP address
Priority interface state tracking	Disabled
Advertisement interval	1 second for IPv4 100 centiseconds for IPv6

Configuring IP Services

This section includes the following topics:

- [Configuring Management Interface, page 5-9](#)
- [Configuring the Default Gateway, page 5-10](#)
- [Configuring Default Networks using IPV4, page 5-11](#)
- [Configuring an IPv4 Address in a VSAN, page 5-11](#)
- [Enabling IPv4 Routing, page 5-11](#)
- [Configuring IPv4 Static Routes, page 5-12](#)
- [Configuring Overlay VSANs, page 5-12](#)
- [Configuring DNS Server, page 5-24](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Management Interface

Detailed Steps

To configure the mgmt0 Ethernet interface for IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface mgmt0 switch(config-if)#	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	switch(config-if)# ip address 10.1.1.1 255.255.255.0	Enters the IPv4 address (10.1.1.1) and IPv4 subnet mask (255.255.255.0) for the management interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

To configure the mgmt0 Ethernet interface for IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface mgmt0 switch(config-if)#	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	switch(config-if)# ipv6 address 2001:0db8:800:200c::417a/64	Enters the IPv6 address (2001:0DB8:800:200C::417A) and IPv6 prefix length (/64) for the management interface and enables IPv6 processing on the interface.
	switch(config-if)# ipv6 enable	Automatically configures a link-local IPv6 address on the interface and enables IPv6 processing on the interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

To configure the mgmt0 Ethernet interface using Device Manager for IPv6, follow these steps:

-
- Step 1 [Select Interface > Mgmt > Mgmt0.](#)
 - Step 2 [Enter the description.](#)
 - Step 3 [Select the administrative state of the interface.](#)
 - Step 4 [Check the CDP check box to enable CDP.](#)
 - Step 5 [Enter the IP address mask.](#)
 - Step 6 [Click Apply to apply the changes.](#)
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring the Default Gateway

Detailed Steps

To configure the default gateway, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ip default-gateway 1.12.11.1	Configures the IPv4 address for the default gateway.

To configure an IP route, follow these steps:

-
- Step 1** Select **Switches > Interfaces > Management**, and select **IP** in the Physical Attributes pane.
- Step 2** Click the **Route** tab in the information pane.
- You see the IP route window showing the switch name, destination, mask, gateway, metric, interface, and active status of each IP route.
- Step 3** Click the **Create Row** icon to add a new IP route.
- Step 4** Complete the fields in this window.
- Enter the switch name in the Switch field.
 - Configure a static route, by entering the destination network ID and subnet mask in the Routedest and Mask fields.
 - Configure a default gateway by entering the IP address of the seed switch in the Gateway field.
 - Set the Metric and Interface fields.



Note

With Cisco NX-OS Release 4.2(1) and later, CPP interfaces also are available for selection when you create a new IP route.

-
- Step 5** Click the **Create** icon.

To configure an IP route or identify the default gateway using Device Manager, follow these steps:

-
- Step 1** Choose **IP > Routes**.
- You see the IP Routes window.
- Step 2** Create a new IP route or identify the default gateway on a switch by clicking **Create**.
- Step 3** Complete the fields in this window.
- Enter the switch name in the Switch field.
 - Configure a static route, by entering the destination network ID and subnet mask in the Routedest and Mask fields.
 - Configure a default gateway by entering the IP address of the seed switch in the Gateway field.
 - Set the Metric and Interface fields.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

With Cisco NX-OS Release 4.2(1) and later, CPP interfaces also are available for selection when you create a new IP route.

If you choose the CPP interface, the switch uses the input CPP-assigned IP address and mask to generate the IP route prefix.

Step 4

Click **Create** to add the IP route.

**Note**

You cannot delete the switch-generated IP route for the CPP interface. If you try to delete the IP route for the CPP interface, SNMP displays this error message:

```
ip: route type not supported.
```

Configuring Default Networks using IPv4

Detailed Steps

To configure default networks using IPv4 addresses, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip default-network 190.10.1.0	Configures the IPv4 address for the default network (190.10.1.0).
	switch(config)# ip route 10.0.0.0 255.0.0.0 131.108.3.4 switch(config)# ip default-network 10.0.0.0	Defines a static route to network 10.0.0.0 as the static default route.

Configuring an IPv4 Address in a VSAN

Detailed Steps

To create a VSAN interface and configure an IPv4 address for that interface, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures the interface for the specified VSAN (10).
Step 3	switch(config-if)# ip address 10.0.0.12 255.255.255.0	Configures the IPv4 address and netmask for the selected interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

Enabling IPv4 Routing

By default, the IPv4 routing feature is disabled in all switches.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To enable the IPv4 routing feature, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip routing	Enables IPv4 routing (disabled by default).
Step 3	switch(config)# no ip routing	Disables IPv4 routing and reverts to the factory settings.

Configuring IPv4 Static Routes

Detailed Steps

To configure an IPv4 static route, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip route <i>network IP address netmask next hop IPv4 address distance number interface vsan number</i> For example: switch(config)# ip route 10.0.0.0 255.0.0.0 20.20.20.10 distance 10 interface vsan 1 switch(config)#	Configures the static route for the specified IPv4 address, subnet mask, next hop, distance, and interface.

Configuring Overlay VSANs

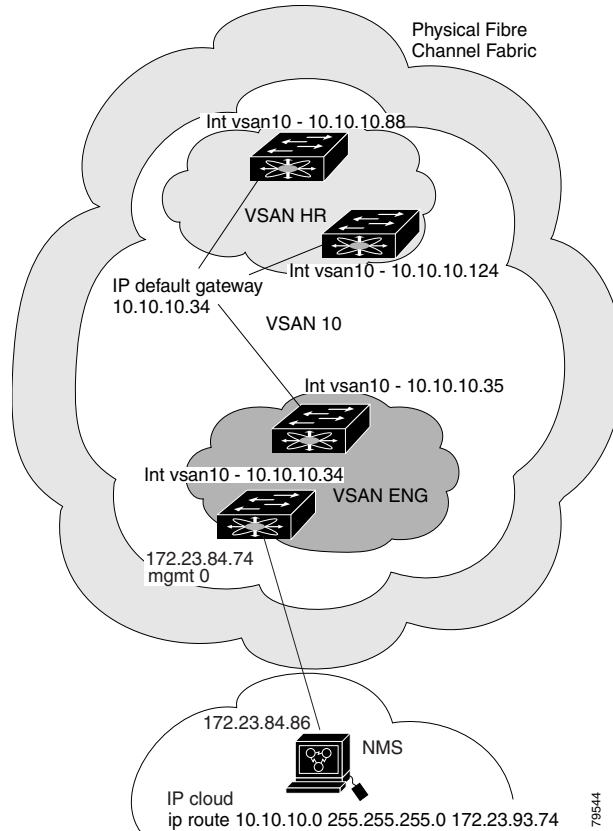
Detailed Steps

To configure an overlay VSAN, follow these steps:

- Step 1 Add the VSAN to the VSAN database on all switches in the fabric.
- Step 2 Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side.
- Step 3 Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.
- Step 4 Configure the default gateway (route) and the IPv4 address on switches that point to the NMS as shown in [Figure 5-5](#).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 5-5 Overlay VSAN Configuration Example



Note

To configure the management interface displayed in Figure 5-5, set the default gateway to an IPv4 address on the Ethernet network.

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in Figure 5-5), follow these steps:

	Command	Purpose
Step 1	switch# confi g t	Enters configuration mode.
Step 2	switch(config)# vsan database switch-config-vsan-db#	Configures the VSAN database.
Step 3	switch--config-vsan-db# vsan 10 name MGMT_VSAN	Defines the VSAN in the VSAN database on all of the switches in the Fibre Channel fabric.
Step 4	switch--config-vsan-db# exit switch(config)#	Exits the VSAN database mode.
Step 5	switch(config)# interface vsan 10 switch(config-if)#	Creates a VSAN interface (VSAN 10).
Step 6	switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0	Assigns an IPv4 address and subnet mask for this switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 7	switch(config-if)# no shutdown	Enables the configured interface.
Step 8	switch(config-if)# end switch#	Exits to EXEC mode.
Step 9	switch# exit	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.

To configure the NMS station displayed in Figure 5-5, follow this step:

	Command	Purpose
Step 1	nms# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.

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.

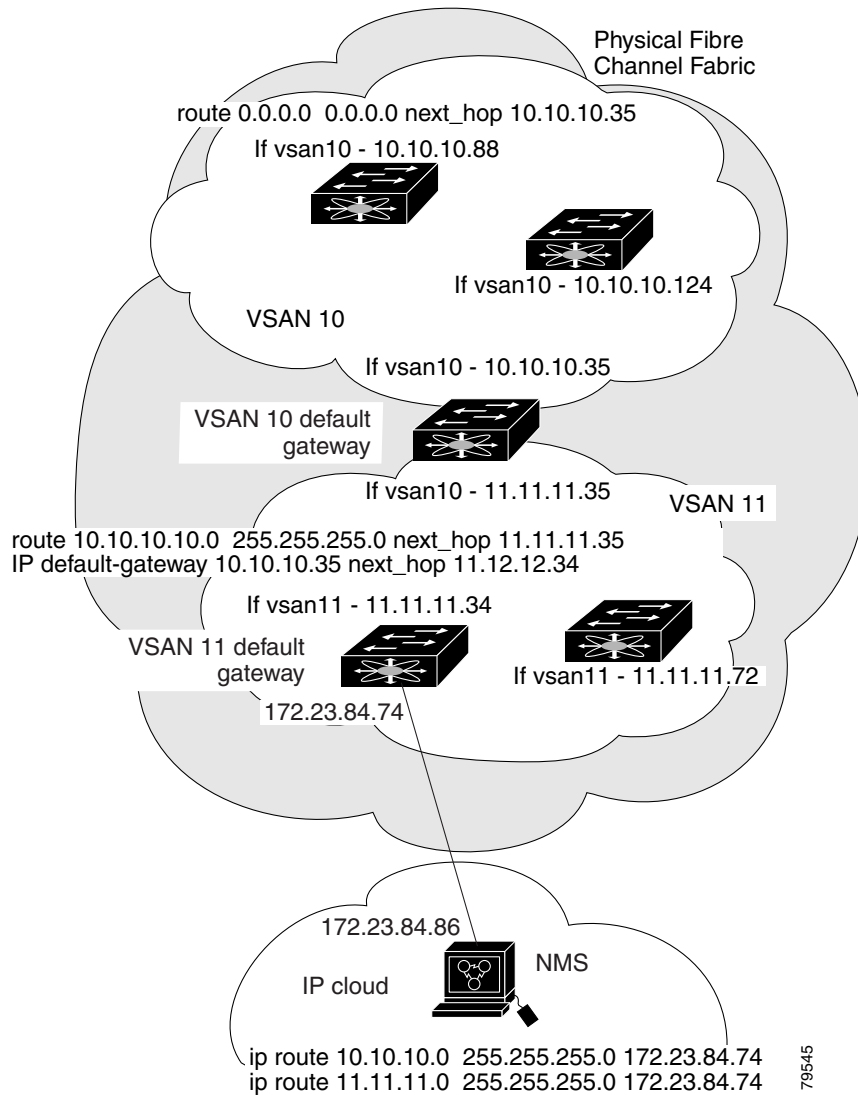
Detailed Steps

To configure multiple VSANs, follow these steps:

- Step 1** Add the VSAN to the VSAN database on any switch in the fabric.
- Step 2** Create a VSAN interface for the appropriate VSAN on any switch in the fabric.
- Step 3** Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.
- Step 4** Define the multiple static routes on the Fibre Channel switches and the IP cloud as shown in Figure 5-6.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 5-6 Multiple VSAN Configuration Example



To configure an overlay VSAN (using the example in Figure 5-6), follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# vsan database switch-config-vsan-db#	Configures the VSAN database.
Step 3	switch-config-vsan-db# vsan 10 name MGMT_VSAN_10 switch-config-vsan-db#	Defines the VSAN in the VSAN database on all of the switches in VSAN 10.
Step 4	switch-config-vsan-db# exit switch(config)#	Exits the VSAN database configuration submenu.
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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 6	switch-config-vsdb# exit switch(config)#	Exits the VSAN database configuration submode.
Step 7	switch(config)# interface vsan 10 switch(config-if)#	Enters the interface configuration submode for VSAN 10.
Step 8	switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0 switch(config-if)#	Assigns an IPv4 address and subnet mask for this interface.
Step 9	switch(config-if)# no shutdown	Enables the configured interface for VSAN 10.
Step 10	switch(config-if)# exit switch(config)#	Exits the VSAN 10 interface mode.
Step 11	switch(config)# interface vsan 11 switch(config-if)#	Enters the interface configuration submode for VSAN 11.
Step 12	switch(config-if)# ip address 11.11.11.0 netmask 255.255.255.0 switch(config-if)#	Assigns an IPv4 address and subnet mask for this interface.
Step 13	switch(config-if)# no shutdown	Enables the configured interface for VSAN 11.
Step 14	switch(config-if)# end switch#	Exits to EXEC mode.
Step 15	switch# exit	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.
Step 16	NMS# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the IPv4 cloud.
Step 17	NMS# route ADD 11.11.11.0 MASK 255.255.255.0 172.22.93.74	Defines a static route for VSAN 11 on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.
Step 18	switch# route 10.10.10.0 255.255.255.0 next_hop 11.11.11.35	Defines the route to reach subnet 10 from subnet 11.

Configuring VRRP

This section describes how to configure VRRP and includes the following topics:

- Adding and Deleting a Virtual Router, page 5-17
- Virtual Router Initiation, page 5-17
- [Adding Virtual Router IP Addresses, page 5-18](#)
- Setting the Priority for the Virtual Router, page 5-19
- Setting the Time Interval for Advertisement Packets, page 5-20

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Configuring or Enabling Priority Preemption, page 5-21](#)
- [Setting Virtual Router Authentication, page 5-22](#)
- [Tracking the Interface Priority, page 5-23](#)

Adding and Deleting a Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.



Note

The total number of VRRP groups that you can configure on a Gigabit Ethernet port, including main interfaces and subinterfaces, cannot exceed seven. This limitation applies to both IPv4 and IPv6 groups.

Detailed Steps

To create or remove a VR for IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates VR ID 250.
	switch(config-if)# no vrrp 250	Removes VR ID 250.

To create or remove a VR for IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp ipv6 250 switch(config-if-vrrp-ipv6)#	Creates VR ID 250.
	switch(config-if)# no vrrp ipv6 250	Removes VR ID 250.

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address, either IPv4 or IPv6, before attempting to enable a VR.

To enable or disable a virtual router configure for IPv4, follow these steps:

	Command	Purpose
Step 1	switch(config-if-vrrp)# no shutdown	Enables VRRP configuration.
	switch(config-if-vrrp)# shutdown	Disables VRRP configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To enable or disable a virtual router configured for IPv6, follow these steps:

	Command	Purpose
Step 1	switch(config-if-vrrp-ipv6) # no shutdown	Enables VRRP configuration.
	switch(config-if-vrrp-ipv6) # shutdown	Disables VRRP configuration.

Adding Virtual Router IP Addresses

One virtual router IP address can be configured for a virtual router. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address. You can configure either an IPv4 address or an IPv6 address.

According to the VRRP specification, the master VRRP router drops the packets addressed to the virtual router's IP address because the virtual router is only intended as a next-hop router to forward packets. In MDS switches however, some applications require that packets addressed to virtual router's IP address be accepted and delivered to them. By using the **secondary** option to the virtual router IPv4 address, the VRRP router will accept these packets when it is the master.

Detailed Steps

To manage IP addresses for virtual routers from Device Manager, follow these steps:

-
- Step 1 Choose **IP > VRRP**. You see the **Operations** tab of the VRRP dialog box.
 - Step 2 Click the **IP Addresses** tab on the VRRP dialog box.
 - Step 3 To create a new VRRP entry, click **Create**. You see the Create VRRP IP Addresses window.
 - Step 4 Complete the fields in this window to create a new VRRP IP address, and click **OK** or **Apply**.
-

To configure an IPv4 address for a virtual router, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# interface ip address 10.0.0.12 255.255.255.0	Configures an IPv4 address and subnet mask. The IPv4 address must be configured before the VRRP is added.
Step 4	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates VR ID 250.
Step 5	switch(config-if-vrrp)# address 10.0.0.10	Configures the IPv4 address for the selected VR.
	switch(config-if-vrrp)# no address 10.0.0.10	Removes the IP address for the selected VR.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 6	switch(config-if-vrrp)# address 10.0.0.10 secondary	Configures the IP address (10.0.0.10) as secondary for the selected VR. Note The secondary option should be used only with applications that require VRRP routers to accept the packets sent to the virtual router's IP address and deliver to them. For example, iSNS requires this option (see the Enabling the iSNS Server).
	switch(config-if-vrrp)# no address 10.0.0.10 secondary	Removes the IP address (10.0.0.10) as secondary for the selected VR.

To configure an IPv6 address for a virtual router, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 12 switch(config-if)#	Configures a VSAN interface (VSAN 12).
Step 3	switch(config-if)# interface ipv6 address 2001:0db8:800:200c::417a/64	Configures an IP address and prefix. The IPv6 address must be configured before the VRRP is added.
Step 4	switch(config-if)# vrrp ipv6 200 switch(config-if-vrrp-ipv6)#	Creates VR ID 200.
Step 5	switch(config-if-vrrp-ipv6)# address 2001:0db8:800:200c::417a	Assigns single primary link-local IPv6 address or one of the multiple secondary IPv6 addresses. Note If this IPv6 address is the same as the physical IPv6 address, this switch is automatically the owner of this IPv6 address.
	switch(config-if-vrrp-ipv6)# no address 2001:0db8:800:200c::417a	Removes the IPv6 address for the selected VR.

Setting the Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for switches with the primary IP address.

Detailed Steps

To set the priority for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	switch(config-if-vrrp)# priority 2	Configures the priority for the selected VRRP. Note Priority 255 cannot be preempted.
	switch(config-if-vrrp)# no priority	Reverts to the default value (100 for switch with the secondary IPv4 addresses and 255 for switches with the primary IPv4 address).

To set the priority for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 12 switch(config-if)#	Configures a VSAN interface (VSAN 12).
Step 3	switch(config-if)# vrrp ipv6 200 switch(config-if-vrrp-ipv6)#	Creates a virtual router.
Step 4	switch(config-if-vrrp-ipv6)# priority 2	Configures the priority for the selected VRRP. Note Priority 255 cannot be preempted.
	switch(config-if-vrrp-ipv6)# no priority	Reverts to the default value (100 for switch with the secondary IPv6 addresses and 255 for switches with the primary IPv6 address).

Setting the Time Interval for Advertisement Packets

The valid time range for an advertisement packet on an interface using IPv4 is between 1 and 255 seconds. The default value is 1 (one) second. If the switch has the primary IP address, this time must be specified.

Detailed Steps

To set the time interval for advertisement packets for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp 50 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# advertisement-interval 15	Sets the interval time in seconds between sending advertisement frames. The range is 1 to 255.
	switch(config-if-vrrp)# no advertisement-interval	Reverts to the default value (1 second).

Send documentation comments to dcnm-san-docfeedback@cisco.com

To set the time interval for advertisement packets for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 12 switch(config-if)#	Configures a VSAN interface (VSAN 12).
Step 3	switch(config-if)# vrrp ipv6 200 switch(config-if-vrrp-ipv6)#	Creates a virtual router.
Step 4	switch(config-if-vrrp-ipv6)# advertisement-interval 150	Sets the interval time in centiseconds between sending advertisement frames. The range is 100 to 4095. The default is 100 centiseconds.
	switch(config-if-vrrp-ipv6)# no advertisement-interval	Reverts to the default value (100 centiseconds).

Configuring or Enabling Priority Preemption

You can enable a higher-priority backup virtual router to preempt the lower-priority master virtual router.



Note

If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.



Note

The VRRP preemption is not supported on IP storage Gigabit Ethernet interfaces.

Detailed Steps

To enable or disable preempting when using IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# preempt	Enables the higher priority backup virtual router to preempt the lower priority master virtual router. Note This preemption does not apply to the primary IP address.
	switch(config-if-vrrp)# no preempt	Disables (default) the preempt option and allows the master to keep its priority level.

To enable or disable preempting when using IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 12 switch(config-if)#	Configures a VSAN interface (VSAN 12).
Step 3	switch(config-if)# vrrp ipv6 200 switch(config-if-vrrp-ipv6)#	Creates a virtual router.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 4	<code>switch(config-if-vrrp-ipv6)# preempt</code>	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.
	<code>switch(config-if-vrrp-ipv6)# no preempt</code>	Disables (default) the preempt option and allows the master to keep its priority level.

Setting Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.



Note All VRRP configurations must be duplicated.



Note VRRP router authentication does not apply to IPv6.

Detailed Steps

To set an authentication option for a virtual router, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# interface vsan 1</code> <code>switch(config-if)#</code>	Configures a VSAN interface (VSAN 1).
Step 3	<code>switch(config-if)# vrrp 250</code> <code>switch(config-if-vrrp)#</code>	Creates a virtual router.
Step 4	<code>switch(config-if-vrrp)# authentication text password</code>	Assigns the simple text authentication option and specifies the password for this option.
	<code>switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003</code>	Assigns the MD5 authentication option and specifies the key and the unique SPI value for this option. The SPI and the valid range is 0x100 to 0xFFFFFFFF.
	<code>switch(config-if-vrrp)# no authentication</code>	Assigns the no authentication option, which is the default.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Tracking the Interface Priority

Interface state tracking changes the priority of the virtual router based on the state of another interface in the switch. When the tracked interface is down, the priority reverts to the priority value for the virtual router (see the “[Setting the Priority for the Virtual Router](#)” section on page 5-19). When the tracked interface is up, the priority of the virtual router is restored to the interface state tracking value. You can track the state of either a specified VSAN interface or the management interface (mgmt 0). The interface state tracking feature is disabled by default.



Note

For interface state tracking to function, you must enable preemption on the interface.

Detailed Steps

To track the interface priority for a virtual router using IPv4, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 10 switch(config-if)#	Configures a VSAN interface (VSAN 10).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# preempt	Enables priority preemption.
Step 5	switch(config-if-vrrp)# track interface mgmt 0 priority 2	Specifies the priority of the virtual router to be modified based on the state of the management interface.
	switch(config-if-vrrp)# no track	Disables the tracking feature.

To track the interface priority for a virtual router using IPv6, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 12 switch(config-if)#	Configures a VSAN interface (VSAN 12).
Step 3	switch(config-if)# vrrp ipv6 200 switch(config-if-vrrp-ipv6)#	Creates a virtual router.
Step 4	switch(config-if-vrrp-ipv6)# preempt	Enables priority preemption.
Step 5	switch(config-if-vrrp-ipv6)# track interface mgmt 0 priority 2	Specifies the priority of the virtual router to be modified based on the state of the management interface.
	switch(config-if-vrrp-ipv6)# no track	Disables the tracking feature.

Note You must enable IPv6 on the tracked interface for the priority tracking to take affect (see the “[Configuring Basic Connectivity for IPv6](#)” section on page 8-13). If IPv6 is not enabled, the interface state is treated as down by VRRP over IPv6, regardless of the actual state of the interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring DNS Server

Detailed Steps

To configure a DNS server, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ip domain-lookup	Enables the IP Domain Naming System (DNS)-based host name-to-address translation.
	switch(config)# no ip domain-lookup	Disables (default) the IP DNS-based host name-to-address translation and reverts to the factory default.
Step 3	switch(config)# ip domain-name cisco.com	Enables the default domain name feature used to complete unqualified host names. Any IP host name that does not contain a domain name (that is, any name without a dot) will have the dot and cisco.com appended to it before being added to the host table.
	switch(config)# no ip domain-name cisco.com	Disables (default) the domain name.
Step 4	switch(config)# ip domain-list harvard.edu switch(config)# ip domain-list stanford.edu switch(config)# ip domain-list yale.edu	Defines a filter of default domain names to complete unqualified host names by using the ip domain-list global configuration command. You can define up to 10 domain names in this filter. To delete a name from a filter, use the no form of this command.
	switch(config)# no ip domain-list	Deletes the defined filter and reverts to factory default. No domains are configured by default.
Note If you have not configured a domain list, the domain name that you specified with the ip domain-name global configuration command is used. If you configured a domain list, the default domain name is not used. The ip domain-list command is similar to the ip domain-name command, except that with the ip domain-list command you can define a list of domains, each to be tried in turn.		
Step 5	switch(config)# ip name-server 15.1.0.1 2001:0db8:800:200c::417a	Specifies the first address (15.1.0.1) as the primary server and the second address (2001:0db8:800:200c::417a) as the secondary server. You can configure a maximum of six servers.
	switch(config)# no ip name-server	Deletes the configured server(s) and reverts to factory default. No server is configured by default.
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.	

Send documentation comments to dcnm-san-docfeedback@cisco.com

Verifying IP Services Configuration

To display IP services configuration information, perform one of the following tasks:

Command	Purpose
<code>show ip routing</code>	Displays the IP routing status.
<code>show arp</code>	Displays the ARP table.
<code>switch(config)# no arp 172.2.0.1</code>	Removes an ARP entry from the ARP table.
<code>clear arp-cache</code>	Delete all entries from the ARP table. The ARP table is empty by default.
<code>show vrrp vr 7 interface vsan 2 configuration</code>	Displays IPv4 VRRP configured information
<code>show vrrp vr 7 interface vsan 2 status</code>	Displays IPv4 VRRP status information.
<code>show vrrp vr 7 interface vsan 2 statistics</code>	Displays IPv4 VRRP statistics.
<code>show vrrp ipv6 vr 1</code>	Displays IPv6 VRRP information.
<code>show vrrp ipv6 vr 1 interface gigabitethernet 1/5 configuration</code>	Displays IPv6 VRRP interface configuration information.
<code>show vrrp ipv6 vr 1 interface gigabitethernet 1/5 status</code>	Displays IPv6 VRRP interface status information.
<code>show vrrp ipv6 vr 1 interface gigabitethernet 1/5 statistics</code>	Displays IPv6 VRRP statistics.
<code>show vrrp statistics</code>	Displays VRRP cumulative statistics.
<code>switch# clear vrrp Statistics</code>	Clears VRRP statistics.
<code>clear vrrp vr 1 interface vsan 1</code>	Clears VRRP statistics on a specified interface.
<code>clear vrrp ipv4 vr 7 interface vsan 2</code>	Clears VRRP IPv4 statistics on a specified interface.
<code>clear vrrp ipv6 vr 7 interface vsan 2</code>	Clears VRRP IPv6 statistics on a specified interface.
<code>show hosts</code>	Displays configured host details.

This section includes the following topics:

- Verifying the Default Gateway Configuration, page 5-25
- Verifying the VSAN Interface Configuration, page 5-26
- Verifying the IPv4 Routing Configuration, page 5-26
- Verifying IPv4 Static Route Information, page 5-26
- Displaying and Clearing ARPs, page 5-27
- Displaying DNS Host Information, page 5-30

Verifying the Default Gateway Configuration

Use the `show ip route` command to verify the default gateway configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch# show ip route

Codes: C - connected, S - static

Gateway of last resort is 1.12.11.1

S 5.5.5.0/24 via 1.1.1.1, GigabitEthernet1/1
C 1.12.11.0/24 is directly connected, mgmt0
C 1.1.1.0/24 is directly connected, GigabitEthernet1/1
C 3.3.3.0/24 is directly connected, GigabitEthernet1/6
C 3.3.3.0/24 is directly connected, GigabitEthernet1/5
S 3.3.3.0/24 via 1.1.1.1, GigabitEthernet1/1
```

Verifying the VSAN Interface Configuration

Use the **show interface vsan** command to verify the configuration of the VSAN interface.



Note

You can see the output for this command only if you have previously configured a VSAN interface.

```
switch# show interface vsan 1
vsan1 is down (Administratively down)
  WWPN is 10:00:00:0c:85:90:3e:85, FCID not assigned
  Internet address is 10.0.0.12/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
```

Verifying the IPv4 Routing Configuration

Use the **show ip routing** command to verify the IPv4 routing configuration.

```
switch(config)# show ip routing
ip routing is enabled
```

Verifying IPv4 Static Route Information

Use the **show ip route** command to verifying the IPv4 static route configuration.

```
switch# show ip route configured
Destination          Gateway             Mask Metric         Interface
-----
          default          172.22.95.1         0.0.0.0     0             mgmt0
          10.1.1.0            0.0.0.0            255.255.255.0 0             vsan1
          172.22.95.0         0.0.0.0            255.255.255.0 0             mgmt0
```

Use the **show ip route** command to verifying the active and connected IPv4 static route.

```
switch# show ip route

Codes: C - connected, S - static

Default gateway is 172.22.95.1
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
C 172.22.95.0/24 is directly connected, mgmt0
C 10.1.1.0/24 is directly connected, vsan1
```

Example 5-1 Displays the IP Routing Status

```
switch# show ip routing
ip routing is disabled
```

Displaying and Clearing ARPs

Address Resolution Protocol (ARP) entries in Cisco MDS 9000 Family switches can be displayed, deleted, or cleared. The ARP feature is enabled on all switches.

- Use the **show arp** command to display the ARP table.

```
switch# show arp
Protocol Address          Age (min)  Hardware Addr  Type   Interface
Internet 171.1.1.1              0  0006.5bec.699c  ARPA  mgmt0
Internet 172.2.0.1              4  0000.0c07.ac01  ARPA  mgmt0
```

- Use the **no arp** command in configuration mode to remove an ARP entry from the ARP table.

```
switch(config)# no arp 172.2.0.1
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
```

Displaying IPv4 VRRP Information

Use the **show vrrp vr** command to display configured IPv4 VRRP information (see Examples 5-2 to 5-4).

Example 5-2 Displays IPv4 VRRP Configured Information

```
switch# show vrrp vr 7 interface vsan 2 configuration
vr id 7 configuration
admin state down
priority 100
no authentication
advertisement-Interval 1
preempt yes
tracking interface vsan1 priority 2
protocol IP
```

Example 5-3 Displays IPv4 VRRP Status Information

```
switch# show vrrp vr 7 interface vsan 2 status
vr id 7 status
MAC address 00:00:5e:00:01:07
Operational state: init
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Example 5-4 Displays IPv4 VRRP Statistics

```
switch# show vrrp vr 7 interface vsan 2 statistics
vr id 7 statistics
Become master 0
Advertisement 0
Advertisement Interval Error 0
Authentication Failure 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Authentication Type 0
Mismatch Authentication 0
Invalid Packet Length 0
```

Displaying IPv6 VRRP Information

Use the `show vrrp ipv6 vr` command to display configured IPv6 VRRP information (see Example 5-5 through Example 5-9).

Example 5-5 Displays IPv6 VRRP Information

```
switch# show vrrp ipv6 vr 1
-----
Interface  VR IpVersion Pri   Time Pre State  VR IP addr
-----
GigE1/5   1   IPv6    100 100cs  master 2004::1
GigE1/6   1   IPv6    100 100cs  backup 2004::1
```

Example 5-6 Displays IPv6 VRRP Interface Configuration Information

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 configuration
IPv6 vr id 1 configuration
admin state up
priority 100
associated ip: 2004::1
advertisement-interval 100
preempt no
protocol IPv6
```

Example 5-7 Displays IPv6 VRRP Interface Status Information

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 status
IPv6 vr id 1 status
MAC address 00:00:5e:00:02:01
Operational state: master
Up time 37 min, 10 sec
Master IP address: fe80::20c:30ff:fedc:96dc
```

Example 5-8 Displays IPv6 VRRP Statistics

```
switch# show vrrp ipv6 vr 1 interface gigabitethernet 1/5 statistics
IPv6 vr id 1 statistics
Become master 1
Advertisement 0
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
Advertisement Interval Error 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Packet Length 0
```

Displaying VRRP Statistics

Use the **show vrrp statistics** command to display configured IPv6 VRRP information (see Example 5-9).

Example 5-9 Displays VRRP Cumulative Statistics

```
switch# show vrrp statistics
Invalid checksum 0
Invalid version 0
Invalid VR ID 0
```

Clearing VRRP Statistics

Use the **clear vrrp statistics** command to clear all the VRRP statistics for all interfaces on the switch (see Example 5-10).

Example 5-10 Clears VRRP Statistics

```
switch# clear vrrp Statistics
```

Use the **clear vrrp vr** command to clear both the IPv4 and IPv6 VRRP statistics for a specified interface (see Example 5-10).

Example 5-11 Clears VRRP Statistics on a Specified Interface

```
switch# clear vrrp vr 1 interface vsan 1
```

Use the **clear vrrp ipv4** command to clear all the statistics for the specified IPv4 virtual router (see Example 5-12).

Example 5-12 Clears VRRP IPv4 Statistics on a Specified Interface

```
switch# clear vrrp ipv4 vr 7 interface vsan 2
```

Use the **clear vrrp ipv6** command to clear all the statistics for the specified IPv6 virtual router (see Example 5-13).

Example 5-13 Clears VRRP IPv6 Statistics on a Specified Interface

```
switch# clear vrrp ipv6 vr 7 interface vsan 2
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see Example 5-14).

Example 5-14 Displays Configured Host Details

```
switch# show hosts
Default domain is cisco.com
Domain list: ucsc.edu harvard.edu yale.edu stanford.edu
Name/address lookup uses domain service
Name servers are 15.1.0.1 15.2.0.0
```


Configuration Examples for IP Services

This section describe an example configuration for IPFC. Figure 5-7 shows an example network.

The example network has the following links:

- Switch_1 is connected to the main network by the mgmt 0 interface and to the fabric by an ISL.
- Switch_2 and Switch_3 are connected to the fabric by an ISL but are not connected to the main network.

Figure 5-7 IPFC Example Network



The following steps show how to configure Switch_1 in the example network in Figure 5-7:

Step 1 Create the VSAN interface and enter interface configuration submode.

```
switch_1# config t
switch_1(config)# interface vsan 1
switch_1(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_1(config-if)# ip address 10.1.1.1 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_1(config-if)# no shutdown
```


Send documentation comments to dcnm-san-docfeedback@cisco.com

```
switch_1(config-if)# exit
switch_1(config)#
```

Step 4 Enable IPv4 routing.

```
switch_1(config)# ip routing
switch_1(config)# exit
switch_1#
```

Step 5 Display the routes.

```
switch_1# show ip route

Codes: C - connected, S - static

C 172.16.1.0/23 is directly connect, mgmt0
C 10.0.0.0./8 is directly connected, vsan1
```

The following steps show how to configure Switch_2 in the example network in Figure 5-7:

Step 1 Enable the mgmt 0 interface.



Note Configure this switch using the console connection.

```
switch_2# config t
switch_2(config)# interface mgmt 0
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 2 Create the VSAN interface and enter interface configuration submode.

```
switch_2# config t
switch_2(config)# interface vsan 1
switch_2(config-if)#
```

Step 3 Configure the IP address and subnet mask.

```
switch_2(config-if)# ip address 10.1.1.2 255.0.0.0
```

Step 4 Enable the VSAN interface and exit interface configuration submode.

```
switch_2(config-if)# no shutdown
switch_2(config-if)# exit
switch_2(config)#
```

Step 5 Enable IPv4 routing.

```
switch_2(config)# ip routing
switch_2(config)# exit
switch_2#
```

Step 6 Display the routes.

```
switch_2# show ip route

Codes: C - connected, S - static

C 10.0.0.0./8 is directly connected, vsan1
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 7 Verify the connectivity to Switch_1.

```
switch_2# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.618 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.528 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.567 ms

--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 4998 ms
rtt min/avg/max/mdev = 0.528/0.570/0.618/0.057 ms
```

The following steps show how to configure Switch_3 in the example network in Figure 5-7:

Step 1 Enable the mgmt 0 interface.



Note Configure this switch using the console connection.

```
switch_3# config t
switch_3(config)# interface mgmt 0
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#
```

```
switch_3# config t
switch_3(config)# interface vsan 1
switch_3(config-if)#
```

Step 2 Configure the IP address and subnet mask.

```
switch_3(config-if)# ip address 10.1.1.3 255.0.0.0
```

Step 3 Enable the VSAN interface and exit interface configuration submode.

```
switch_3(config-if)# no shutdown
switch_3(config-if)# exit
switch_3(config)#
```

Step 4 Enable IPv4 routing.

```
switch_3(config)# ip routing
switch_3(config)# exit
switch_3#
```

Step 5 Display the routes.

```
switch_3# show ip route

Codes: C - connected, S - static

C 10.0.0.0/8 is directly connected, vsan1
```

Step 6 Verify the connectivity to Switch_1.

```
switch_3# ping 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=1.19 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.510 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.653 ms
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
--- 10.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2008 ms
rtt min/avg/max/mdev = 0.510/0.787/1.199/0.297 ms
```

Field Descriptions for IP Services

This section describes the field descriptions.

IP Routes

Field	Description
Routing Enabled	When this check box is enabled, the switch is acting as in IP router.
Destination, Mask, Gateway	The value that identifies the local interface through which the next hop of this route should be reached.
Metric	The primary routing metric for this route.
Interface	The local interface through which the next hop of this route should be reached.
Active	Indicates whether the route is active.

IP Statistics ICMP

Field	Description
InParmProbs	The number of ICMP Parameter Problem messages received.
OutParmProbs	The number of ICMP Parameter Problem messages sent.
InSrcQuenchs	The number of ICMP Source Quench messages received.
InRedirects	The number of ICMP Redirect messages received.
InEchoes	The number of ICMP Echo (request) messages received.
InEchoReps	The number of ICMP Echo Reply messages received.
InTimestamps	The number of ICMP Timestamp (request) messages received.
InTimestampReps	The number of ICMP Timestamp Reply messages received.
InAddrMasks	The number of ICMP Address Mask Request messages received.
InAddrMaskReps	The number of ICMP Address Mask Reply messages received.
InDestUnreachs	The number of ICMP Destination Unreachable messages received.
InTimeExcds	The number of ICMP Time Exceeded messages received.
OutSrcQuenchs	The number of ICMP Source Quench messages sent.
OutRedirects	The number of ICMP Redirect messages sent. For a host, this value will always be N/A, since hosts do not send redirects.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
OutEchos	The number of ICMP Echo (request) messages sent.
OutEchoReps	The number of ICMP Echo Reply messages sent.
OutTimestamps	The number of ICMP Timestamp (request) messages sent.
OutTimestampReps	The number of ICMP Timestamp Reply messages sent.
OutAddrMasks	The number of ICMP Address Mask Request messages sent.
OutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.
OutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
OutTimeExcds	The number of ICMP Time Exceeded messages sent.

IP Statistics IP

Field	Description
InHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
InAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. For entities that are not IP routers, and do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
InUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such frames met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any frames counted in ipForwDatagrams which meet this no-route criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down.
FragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example, because their Don't Fragment flag was set.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
ReasmFails	The number of failures detected by the IP reassembly algorithm (for example, timed out, errors, etc). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
InReceives	The total number of input datagrams received from interfaces, including those received in error.
InDelivers	The total number of input datagrams successfully delivered to IP user protocols (including ICMP).
OutRequests	The total number of IP datagrams which local IP user protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP routers, this counter will include only those frames which were source-routed via this entity, and the Source-Route option processing was successful.
FragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
FragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.

IP Statistics SNMP

Field	Description
BadVersions	The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
BadCommunityNames	The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
BadCommunityUses	The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.
ASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
TooBigs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
SilentDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
ProxyDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in a manner (other than a timeout) such that no Response-PDU could be returned.
NoSuchNames	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
BadValues	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
ReadOnlys	The total number valid SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It should be noted that it is a protocol error to generate an SNMP PDU which contains the value readOnly in the error-status field, as such this is provided as a means of detecting incorrect implementations of the SNMP.
GenErrs	The total number of SNMP PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
Pkts	The total number of messages delivered to the SNMP entity from the transport service.
GetRequests	The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol entity.
GetNexts	The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol entity.
SetRequests	The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol entity.
OutTraps	The total number of SNMP Trap PDUs which have been generated by the SNMP protocol entity.
OutGetResponses	The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol entity.
OutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
TotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
TotalSetVars	The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IP Statistics UDP

Field	Description
InErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
InDatagrams	The total number of UDP datagrams delivered to UDP users.
OutDatagrams	The total number of UDP datagrams sent from this entity.
NoPorts	The total number of received UDP datagrams for which there was no application at the destination port.

mgmt0 Statistics

Field	Description
InErrors	Total number of received errors on the interface.
OutErrors	Total number of transmitted errors on the interface.
InDiscards	Total number of received discards on the interface.
OutDiscards	Total number of transmitted discards on the interface.
RxBytes	Total number of bytes received.
TxBytes	Total number of bytes transmitted.
RxFrames	Total number of frames received.
TxFrames	Total number of frames transmitted.

TCP UDP TCP

Field	Description
State	The state of this TCP connection.

TCP UDP UDP

Field	Description
Port	The local port number for this UDP listener.

Send documentation comments to dcnm-san-docfeedback@cisco.com

VRRP General

Field	Description
IP Address Type, Vrid, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
Admin	The admin state of the virtual router (active or notInService).
Oper	The current state of the virtual router. There are three defined values: <ul style="list-style-type: none"> initialize— Indicates that all the virtual router is waiting for a startup event. backup— Indicates the virtual router is monitoring the availability of the master router. master— Indicates that the virtual router is forwarding frames for IP addresses that are associated with this router.
Priority	Specifies the priority to be used for the virtual router master election process. Higher values imply higher priority. A priority of 0 is sent by the master router to indicate that this router has ceased to participate in VRRP and a backup virtual router should transition to become a new master. A priority of 255 is used for the router that owns the associated IP address(es).
AdvInterval	The time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
PreemptMode	Controls whether a higher priority virtual router will preempt a lower priority master.
UpTime	When this virtual router transitioned out of initialized.
Version	The VRRP version on which this VRRP instance is running.
AcceptMode	Controls whether a virtual router in Master state will accept packets addressed to the address owner's IPv6 address as its own if it is not the IPv6 address owner. If true, the virtual router in Master state will accept. If false, the virtual router in Master state will not accept.

VRRP IP Addresses

Field	Description
Interface, VRRP ID, IP Address	Interface, Virtual Router Redundancy Protocol ID, and associated IP address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

VRRP Statistics

Field	Description
IP Address Type, Vrid, Interface	The IP address type (IPv4, IPv6, or DNS), the virtual router ID, and the interface.
LastAdvRx	The total number of VRRP advertisements received by this virtual router.
Protocol Traffic MasterIpAddr	The master router's real (primary) IP address. This is the IP address listed as the source in VRRP advertisement last received by this virtual router.
Protocol Traffic BecomeMaster	The total number of times that this virtual router's state has transitioned to MASTER.
Priority 0 Rx	The total number of VRRP frames received by the virtual router with a priority of 0.
Priority 0Tx	The total number of VRRP frames sent by the virtual router with a priority of 0.
AuthErrors InvalidType	The total number of frames received with an unknown authentication type.
Other Errors dvIntervalErrors	The total number of VRRP advertisement frames received for which the advertisement interval is different than the one configured for the local virtual router.
Other Errors IpTtlErrors	The total number of VRRP frames received by the virtual router with IP TTL (time-to-live) not equal to 255.
Other Errors InvalidTypePktsRcvd	The number of VRRP frames received by the virtual router with an invalid value in the type field.
Other Errors AddressListErrors	The total number of frames received for which the address list does not match the locally configured list for the virtual router.
OtherErrors PacketLengthErrs	The total number of frames received with a frame length less than the length of the VRRP header.
RefreshRate	The interval of time between refreshes.

CDP General

Field	Description
Enable	Whether the Cisco Discovery Protocol is currently running. Entries in CacheTable are deleted when CDP is disabled.
MessageInterval	The interval at which CDP messages are to be generated. The default value is 60 seconds.
HoldTime	The time for the receiving device holds CDP message. The default value is 180 seconds.
LastChange	When the cache table was last changed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CDP Neighbors

Field	Description
Switch	The Internet address for this entity.
Local Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
DeviceName	The remote device's name. By convention, it is the device's fully qualified domain name.
DeviceID	The device ID string as reported in the most recent CDP message.
DevicePlatform	The version string as reported in the most recent CDP message.
Interface	The port ID string as reported in the most recent CDP message.
IPAddress	The (first) network-layer address of the device's SNMP-agent as reported in the address TLV of the most recently received CDP message.
NativeVLAN	The remote device's interface's native VLAN, as reported in the most recent CDP message. The value 0 indicates no native VLAN field (TLV) was reported in the most recent CDP message.
PrimaryMgmtAddr	Indicates the (first) network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.
SecondaryMgmtAddr	Indicates the alternate network layer address at which the device will accept SNMP messages as reported in the most recently received CDP message.

iSNS Profiles

Field	Description
Addr	The address of the iSNS server.
Port	The TCP port of the iSNS server.

iSNS Servers

Field	Description
Name	The name of the iSNS server.
TcpPort	The TCP port used for iSNS messages. If TCP is not supported by this server, the value is 0.
Uptime	The time the server has been active.
ESI Non Response Threshold	The number of ESI messages that will be sent without receiving a response before an entity is unregistered from the iSNS database.
# Entities	The number of entities registered in iSNS on the server.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
# Portals	The number of portals registered in iSNS on the server.
# Portal Groups	The number of portal groups registered in iSNS on the server.
# iSCSI Devices	The number of iSCSI Nodes registered in iSNS on the server.

iSNS Entities

Field	Description
Entity ID	The iSNS entity identifier for the entity.
Last Accessed	The time the entity was last accessed.

iSNS Cloud Discovery

Field	Description
AutoDiscovery	Whether automatic cloud discovery is turned on or off.
DiscoveryDelay	Time duration between successive IP cloud discovery runs.
Discovery	The IP network discovery command to be executed. <ul style="list-style-type: none"> all- Run IP network discovery for all the gigabit Ethernet interfaces in the fabric. noOp (default)- No operation is performed.
CommandStatus	The status of the license install / uninstall / update operation. <ul style="list-style-type: none"> success— Discovery operation completed successfully. nProgress— Discovery operation is in progress. none— No discovery operation is performed. NoIpNetworkNameSpecified— IP Cloud name not specified. invalidNetworkName— IP Cloud is not configured. NoIPSPortNameSpecified— Gigabit Ethernet port if index not specified. invalidIPSPortName— Invalid Gigabit Ethernet port interface. generalISNSFailure— General iSNS server failure.

iSNS Clouds

Field	Description
Id	The ID of the IP cloud.
Switch WWN	The WWN of the switch in this table.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSNS Cloud Interfaces

Field	Description
Name, Switch WWN, Interface, Address	The name, switch WWN, interface, and address of the cloud.

Monitor Dialog Controls

Field	Description
Line Chart	Opens a new window with a line chart representation of the data.
Area Chart	Opens a new window with an area chart representation of the data.
Bar Chart	Opens a new window with a bar chart representation of the data.
Pie Chart	Opens a new window with a pie chart representation of the data.
Reset Cumulative Counters	Resets the counters to 0 if the Column Data display mode is set to Cumulative.
Export to File	Opens a standard Save dialog box. The data is saved as a .TXT file.
Print	Opens a standard Print dialog box.
Update Frequency	The interval at which the data is updated in the monitor dialog.
Column Data	Specifies the type of data that is displayed in the monitor dialog. <ul style="list-style-type: none"> Absolute Value— Displays the total amount since the switch was booted. This is the default for error monitoring. Cumulative—Displays the total amount since the dialog was opened. You can reset the counters by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data. Minimum/sec— Displays the minimum value per second at every refresh interval. Maximum/sec— Displays the maximum value per second at every refresh interval. Last Value/sec— Displays the most recent value per second at every refresh interval. This is the default setting for traffic monitoring.
Elapsed	The amount of time that has elapsed since the dialog was opened. You can reset this counter by clicking the Reset Cumulative Counters button, to gather a new set of cumulative data.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSNS Details iSCSI Nodes

Field	Description
Name	The iSCSI name of the initiator or target associated with the storage node.
Type	The Node Type bit-map defining the functions of this iSCSI node, where 31 is a Target, 30 is an Initiator, 29 is a Control, and all others are reserved.
Alias	The Alias name of the iSCSI node.
ScnBitmap	The State Change Notification (SCN) bitmap for a node.
WWN Token	An optional globally unique 64-bit integer value that can be used to represent the world wide node name of the iSCSI device in a Fibre Channel fabric.
AuthMethod	The iSCSI authentication method enabled for this iSCSI node.

iSNS Details Portals

Field	Description
Addr	The Internet address for this portal.
TcpPort	The port number for this portal.
SymName	The optional Symbolic Name for this portal.
EsiInterval	The Entity Status Inquiry (ESI) Interval for this portal.
TCP ESI	The TCP port number used for ESI monitoring.
TCP Scn	The TCP port used to receive SCN messages from the iSNS server.
SecurityInfo	Security attribute settings for the portal as registered in the Portal Security Bitmap attribute.

Additional References

For additional information related to implementing IP storage, see the following section:

- [Related Document, page 5-44](#)
- [Standards, page 5-44](#)
- [RFCs, page 5-44](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 5-44](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note

FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.

This chapter includes the following topics:

- [Information About IP Storage, page 6-1](#)
- [Licensing Requirements for IP Storage, page 6-10](#)
- [Guidelines and Limitations, page 6-10](#)
- [Default Settings, page 6-11](#)
- [Configuring IP Storage, page 6-11](#)
- [Verifying IP Storage Configuration, page 6-15](#)
- [Field Descriptions for IP Storage, page 6-20](#)
- [Additional References, page 6-28](#)

Information About IP Storage

The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices.

Text Part Number:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target.


The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types of storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2).

Gigabit Ethernet ports in these modules can be configured to support the FCIP protocol, the iSCSI protocol, or both protocols simultaneously:

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 6-1 shows how the IPS module is used in different FCIP scenarios.](#)


Figure 6-1 FCIP Scenarios



- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. [Figure 6-2 depicts the iSCSI scenarios in which the IPS module is used.](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 6-2 *iSCSI Scenarios*



This section contains the following topics:

- [IPS Module Upgrade, page 6-3](#)
- [MPS-14/2 Module Upgrade, page 6-4](#)
- [Supported Hardware, page 6-4](#)
- [Gigabit Ethernet Interfaces for IPv4 Configuration, page 6-4](#)
- [Basic Gigabit Ethernet Configuration, page 6-5](#)
- [IPS Module Core Dumps, page 6-5](#)
- [About VLANs for Gigabit Ethernet, page 6-6](#)
- [Interface Subnet Requirements, page 6-7](#)
- [Verifying Gigabit Ethernet Connectivity, page 6-7](#)
- [Gigabit Ethernet High Availability, page 6-8](#)
- [VRRP for iSCSI and FCIP Services, page 6-8](#)
- [About Ethernet PortChannel Aggregation, page 6-9](#)
- [CDP, page 6-10](#)

IPS Module Upgrade



Caution

A software upgrade is only disruptive for the IPS module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

Send documentation comments to dcnm-san-docfeedback@cisco.com

MPS-14/2 Module Upgrade



Caution

A software upgrade is only partially disruptive for the MPS-14/2 module. The NX-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and two Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one or more of the following hardware:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).



Note

In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel ports and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

Gigabit Ethernet Interfaces for IPv4 Configuration

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can be used to perform only iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.



Note

For information about configuring FCIP, see [Chapter 2, “Configuring FCIP.”](#) For information about configuring iSCSI, see [Chapter 4, “Configuring iSCSI.”](#)

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Note The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.



Note To configure IPv6 on a Gigabit Ethernet interface, see the *Security Configuration Guide, Cisco DCNM for SAN*.




Tip Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration

Figure 6-3 shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Figure 6-3 Gigabit Ethernet IPv4 Configuration Example



Note The port on the Ethernet switch to which the Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should be disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in Cisco IOS or the **set port host** command in Catalyst OS.

IPS Module Core Dumps

IPS core dumps are different from the system's kernel core dumps for other modules. When the IPS module's operating system (OS) unexpectedly resets, it is useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

Use the **show cores** command to list these files.

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module because of its large space requirement. They are copied directly to an external TFTP server.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

Interface Descriptions Configuration

See the *Interfaces Configuration Guide, Cisco DCNM for SAN* for details on configuring the switch port description for any interface.

Beacon Mode Configuration

See the *Interfaces Configuration Guide, Cisco DCNM for SAN* for details on configuring the beacon mode for any interface.

Autonegotiation Configuration

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

MTU Frame Size Configuration

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note

The minimum MTU size is 576 bytes.



Tip

MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

Promiscuous Mode Configuration

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name: *slot-number / port-numberVLAN-ID*.

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 6-1](#)).

Table 6-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	



Note

The configuration requirements in [Table 6-1](#) also apply to Ethernet PortChannels.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.



Note

If the connection fails, verify the following, and ping the IP host again:

- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the `up` state.

Send documentation comments to dcnm-san-docfeedback@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.

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see Figure 6-4).

Figure 6-4 VRRP Scenario



All members of the VRRP group (see Figure 6-4) must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch
- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces



Note

You can configure no more than seven VRRP groups, both IPv4 and IPv6, on a Gigabit Ethernet interface, including the main interface and all subinterfaces.

90855

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Ethernet PortChannel Aggregation

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an Ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that FCIP link.



Note

The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3ad protocol.


Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see Figure 6-5).



Note

PortChannel members must be one of these combinations: ports 1–2, ports 3–4, ports 5–6, or ports 7–8.

Figure 6-5 Ethernet PortChannel Scenario



In Figure 6-5, Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.



Note

PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

Send documentation comments to dcnm-san-docfeedback@cisco.com

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. CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS and MPS-14/2 modules.

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.

See the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.

Licensing Requirements for IP Storage

The following table shows the licensing requirements for this feature:

License	License Description
SAN extension over IP package for IPS-8 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP) SAN extension over IP package for IPS-4 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP_IPS4) 	The following features apply to IPS-8 and IPS-4 modules: <ul style="list-style-type: none"> FCIP FCIP compression FCIP write acceleration FCIP tape read acceleration SAN extension tuner features IVR over FCIP IVR NAT over FCIP Network Stimulator
SAN extension over IP package for MPS-14/2 modules <ul style="list-style-type: none"> (SAN_EXTN_OVER_IP_IPS2) 	The following features apply to the MPS-14/2 module and the fixed Cisco MDS 9216i Switch IP ports: <ul style="list-style-type: none"> FCIP Hardware-based FCIP compression FCIP write acceleration FCIP tape read acceleration SAN extension tuner features IVR over FCIP IVR NAT over FCIP

Guidelines and Limitations



Tip

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established**, **precedence**, and **fragments** options are ignored when you apply IPv4-ACLs (containing these options) to Gigabit Ethernet interfaces.
 - If an IPv4-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B, and an IPv4-ACL specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Default Settings

Table 6-2 lists the default settings for IP storage services parameters.

Table 6-2 Default Gigabit Ethernet Parameters

Parameters	Default
IPS core size	Partial

Configuring IP Storage

This section includes the following topics:

- [Configuring IPS Core Dumps, page 6-11](#)
- [Configuring VRRP for Gigabit Ethernet Interfaces, page 6-12](#)
- [Configuring Ethernet PortChannels, page 6-14](#)

Configuring IPS Core Dumps

Detailed Steps

To configure IPS core dumps on the IPS module, follow these steps:

	Command	Purpose
Step 1	switch# conf terminal switch(config)#	Enters configuration mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 2	switch(config)# ips core dump full ips core dump full' successfully set for module 9	Configures a dump of the full core generation for all IPS modules in the switch.
	switch(config)# no ips core dump full ips core dump partial' successfully set for module 9	Configures a dump of the partial core (default) generation for the IPS module in slot 9.

To configure the Gigabit Ethernet interface for the scenario in Figure 6-3, follow these steps:

-
- Step 1** From Cisco DCNM-SAN, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.
- From Device Manager, right-click the Gigabit Ethernet port that you want to configure and choose **Configure....** You see the Gigabit Ethernet configuration dialog box.
- Step 2** Click the **General** tab in Cisco DCNM-SAN, or click the **GigE** tab in Device Manager to display the general configuration options for the interface.
- Step 3** Set the description and MTU value for the interface. The valid value for the MTU field can be a number in the range from 576 to 9000.
- Step 4** Set **Admin** up or down and check the **CDP** check box if you want this interface to participate in CDP.
- Step 5** Set **IpAddress/Mask** with the IP address and subnet mask for this interface.
- Step 6** From Cisco DCNM-SAN, click the **Apply Changes** icon to save these changes, or click the **Undo Changes** icon to discard changes.
- From Device Manager, click **Apply** to save these changes, or click **Close** to discard changes and close the Gigabit Ethernet configuration dialog box.
-

Configuring VRRP for Gigabit Ethernet Interfaces

Detailed Steps

To configure VRRP for Gigabit Ethernet interfaces using IPv4, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.10 255.255.255.0	Assigns the IPv4 address (10.1.1.10) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp 100 switch(config-if-vrrp)	Creates VR ID 100.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 6	switch(config-if-vrrp)# address 10.1.1.100	Configures the virtual IPv4 address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IPv4 address must be in the same subnet as the IPv4 address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IPv4 address.
Step 7	switch(config-if-vrrp)# priority 10	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	switch(config-if-vrrp)# no shutdown	Enables the VRRP protocol on the selected interface.

To configure VRRP for Gigabit Ethernet interfaces using IPv6, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ipv6 address 2001:0db8:800:200c::417a/64	Assigns the IPv6 address for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp ipv6 100 switch(config-if-vrrp-ipv6)	Creates VR ID 100.
Step 6	switch(config-if-vrrp-ipv6)# address 2001:0db8:800:200c::417a	Assigns single primary link-local IPv6 address or one of the multiple secondary IPv6 addresses. Note If this IPv6 address is the same as the physical IPv6 address, this switch is automatically the owner of this IPv6 address.
Step 7	switch(config-if-vrrp-ipv6)# priority 10	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	switch(config-if-vrrp-ipv6)# no shutdown	Enables the VRRP protocol on the selected interface.



Note

If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual IPv4 IP address is also the IPv4 IP address for the interface, then preemption is implicitly applied.

**Note**

If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

Configuring Ethernet PortChannels

The PortChannel configuration specified in the *Interfaces Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS *Interfaces Configuration Guide* also applies to Ethernet PortChannel configurations.

Detailed Steps

To configure Ethernet PortChannels, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface port-channel 10 switch(config-if)#	Configures the specified PortChannel (10).
Step 3	switch(config-if)# ip address 10.1.1.1 255.255.255.0	Enters the IPv4 address (10.1.1.1) and subnet mask (255.255.255.0) for the PortChannel. Note A PortChannel does not have any members when first created.
Step 4	switch(config-if)# no shutdown	Enables the interface.
Step 5	switch(config)# interface gigabitethernet 9/3 switch(config-if)#	Configures the specified Gigabit Ethernet interface (slot 9, port 3).
Step 6	switch(config-if)# channel-group 10 gigabitethernet 9/3 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up switch(config-if)#	Adds Gigabit Ethernet interfaces 9/3 to channel group 10. If channel group 10 does not exist, it is created. The port is shut down.
Step 7	switch(config-if)# no shutdown	Enables the selected interface.
Step 8	switch(config)# interface gigabitethernet 9/4 switch(config-if)#	Configures the specified Gigabit Ethernet interface (slot 9, port 4).
Step 9	switch(config-if)# channel-group 10 gigabitethernet 9/4 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up	Adds Gigabit Ethernet interfaces 9/4 to channel group 10. The port is shut down.
Step 10	switch(config-if)# no shutdown	Enables the selected interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Verifying IP Storage Configuration

To display IP storage configuration information, perform one of the following tasks:

Command	Purpose
<code>show module</code>	Verifies the status of the module.
<code>show interface gigabitethernet 8/1</code>	Displays the gigabit ethernet interface.
<code>show interface gigabitethernet 4/2.100</code>	Displays the gigabit ethernet subinterface.
<code>show ips stats mac interface gigabitethernet 8/1</code>	Displays ethernet MAC statistics.
<code>show ips stats dma-bridge interface gigabitethernet 7/1</code>	Displays DMA-Bridge statistics.
<code>show ips stats tcp interface gigabitethernet 4/1</code>	Displays TCP statistics.
<code>show ips stats tcp interface gigabitethernet 4/1 detail</code>	Displays Detailed TCP statistics.
<code>show ips stats icmp interface gigabitethernet 2/1</code>	Displays ICMP statistics.

This section includes the following topics:

- [Verifying Module Status, page 6-15](#)
- [Displaying Gigabit Ethernet Interface Statistics, page 6-16](#)
- [Displaying Ethernet MAC Statistics, page 6-17](#)
- [Displaying DMA-Bridge Statistics, page 6-17](#)
- [Displaying TCP Statistics, page 6-18](#)

Verifying 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
----  -
1    0      Caching Services Module   DS-X9560-SMAP       ok
2    8      IP Storage Services Module DS-X9308-SMIP       ok <-----IPS-8 module
4    16     2x1GE IPS, 14x1/2Gbps FC Module DS-X9216i-K9-SUP   ok <-----MPS-14/2 module
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     ha-standby
9    4      IP Storage Services Module DS-X9304-SMIP       ok <-----IPS-4 module

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
----  -
1    2.0(1)     0.201      20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
2    2.0(1)     0.201      20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0
4    2.0(1)     0.201      20:c1:00:05:30:00:07:1e to 20:d0:00:05:30:00:07:1e
5    2.0(1)     0.0        --
6    2.0(1)     0.0        --
9    2.0(1)     0.1        22:01:00:05:30:00:07:1e to 22:04:00:05:30:00:07:1e

Mod          Application Image Description          Application Image Version
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```
-----
1      svc-node1                               1.3 (5M)
1      svc-node2                               1.3 (5M)

Mod  MAC-Address(es)                          Serial-Num
---  -
1    00-05-30-01-49-c2 to 00-05-30-01-4a-46  JAB073907EP
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de  JAB064605a2
4    00-05-30-01-7f-32 to 00-05-30-01-7f-38  JAB081405AM
5    00-05-30-00-2c-4e to 00-05-30-00-2c-52  JAB06350B1M
6    00-05-30-00-19-66 to 00-05-30-00-19-6a  JAB073705GL
9    00-0d-bc-2f-d6-00 to 00-0d-bc-2f-d6-08  JAB080804TN
```

* this terminal session

To verify the status of the module, follow these steps:

-
- Step 1** Select a switch in the Fabric pane.
- Step 2** Open the **Switches** folder and select **Hardware** in the Physical Attributes pane.
- You see the status for all modules in the switch in the Information pane.
-

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface gigabitethernet** command on each switch to verify that the interfaces are up and functioning as desired. See Example 6-1 and Example 6-2.

Example 6-1 Displays the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
GigabitEthernet8/1 is up          <-----The interface is in the up state.
  Hardware is GigabitEthernet, address is 0005.3000.a98e
  Internet address is 10.1.3.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  3343 packets input, 406582 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  8 packets output, 336 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Example 6-2 Displays the Gigabit Ethernet Subinterface

```
switch# show interface gigabitethernet 4/2.100
GigabitEthernet4/2.100 is up
  Hardware is GigabitEthernet, address is 0005.3000.abcb
  Internet address is 10.1.2.100/24
  MTU 1500 bytes
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 packets input, 0 bytes
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

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 6-3.



Note

Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

Example 6-3 Displays Ethernet MAC Statistics

```

switch# show ips stats mac interface gigabitethernet 8/1
Ethernet MAC statistics for port GigabitEthernet8/1
  Hardware Transmit Counters
    237 frame 43564 bytes
    0 collisions, 0 late collisions, 0 excess collisions
    0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
  Hardware Receive Counters
    427916 bytes, 3464 frames, 0 multicasts, 3275 broadcasts
    0 bad, 0 runt, 0 CRC error, 0 length error
    0 code error, 0 align error, 0 oversize error
  Software Counters
    3429 received frames, 237 transmit frames
    0 frames soft queued, 0 current queue, 0 max queue
    0 dropped, 0 low memory

```

Displaying DMA-Bridge Statistics

You can display direct memory access (DMA) device statistics using the **show ips stats dma-bridge interface gigabitethernet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns DMA bridge statistics for that interface. See Example 6-4.



Note

Use the physical interface, not the subinterface, to display DMA-bridge statistics.

Example 6-4 Displays DMA-Bridge Statistics

```

switch# show ips stats dma-bridge interface gigabitethernet 7/1
Dma-bridge ASIC Statistics for port GigabitEthernet7/1
  Hardware Egress Counters
    231117 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
  Hardware Ingress Counters
    218255 Good, 0 protocol error, 0 header checksum error
    0 FC CRC error, 0 iSCSI CRC error, 0 parity error
  Software Egress Counters
    231117 good frames, 0 bad header cksum, 0 bad FIFO SOP
    0 parity error, 0 FC CRC error, 0 timestamp expired error
    0 unregistered port index, 0 unknown internal type

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

0 RDL ok, 0 RDL drop (too big), 0 RDL ttl_1
3656368645 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
Software Ingress Counters
218255 Good frames, 0 header cksum error, 0 FC CRC error
0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
0 out of memory drop, 0 queue full drop
0 RDL ok, 0 RDL drop (too big)
Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]

```

This output shows all Fibre Channel frames that ingress or egress from the Gigabit Ethernet port.

Displaying TCP Statistics

Use the **show ips stats tcp interface gigabitethernet** to display and verify TCP statistics. This command takes the main Ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The **detail** option shows all information maintained by the interface. See Example 6-5 and Example 6-6.

Example 6-5 Displays TCP Statistics

```

switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
Connection Stats
 0 active openings, 3 accepts
 0 failed attempts, 12 reset received, 3 established
Segment stats
163 received, 355 sent, 0 retransmitted
 0 bad segments received, 0 reset sent
TCP Active Connections
Local Address      Remote Address      State      Send-Q  Recv-Q
0.0.0.0:3260       0.0.0.0:0           LISTEN     0       0

```

Example 6-6 Displays Detailed TCP Statistics

```

switch# show ips stats tcp interface gigabitethernet 4/1 detail
TCP Statistics for port GigabitEthernet4/1
TCP send stats
355 segments, 37760 bytes
222 data, 130 ack only packets
3 control (SYN/FIN/RST), 0 probes, 0 window updates
0 segments retransmitted, 0 bytes
0 retransmitted while on ethernet send queue, 0 packets split
0 delayed acks sent
TCP receive stats
163 segments, 114 data packets in sequence, 6512 bytes in sequence
0 predicted ack, 10 predicted data
0 bad checksum, 0 multi/broadcast, 0 bad offset
0 no memory drops, 0 short segments
0 duplicate bytes, 0 duplicate packets
0 partial duplicate bytes, 0 partial duplicate packets
0 out-of-order bytes, 1 out-of-order packets
0 packet after window, 0 bytes after window
0 packets after close
121 acks, 37764 ack bytes, 0 ack toomuch, 4 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
8 window updates, 0 window probe

```


[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

```

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 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
  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 6-7.

Example 6-7 Displays ICMP Statistics

```

switch# show ips stats icmp interface gigabitethernet 2/1
ICMP Statistics for port GigabitEthernet2/1
  0 ICMP messages received
  0 ICMP messages dropped due to errors
ICMP input histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply
ICMP output histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field Descriptions for IP Storage

This section describes the following field descriptions.

FCIP Profiles

Field	Description
IP Address	The Internet address for this entity.
Port	A TCP port other than the FCIP well-known port on which the FCIP entity listens for new TCP connection requests.
SACK	Whether the TCP Selective Acknowledgement Option is enabled to allow the receiver end to acknowledge multiple lost frames in a single ACK, enabling faster recovery.
KeepAlive (s)	The TCP keepalive timeout for all links within this entity.
ReTrans MinTimeout (ms)	The TCP minimum retransmit timeout for all the links on this entity.
ReTrans Max	The maximum number of times that the same item of data will be retransmitted over a TCP connection. If delivery is not acknowledged after this number of retransmissions then the connection is terminated.
Send BufSize (KB)	The aggregate TCP send window for all TCP connections on all links within this entity. This value is used for egress flow control. When the aggregate of the data queued on all connections within this entity reaches this value, the sender is flow controlled.
Bandwidth Max (Kb)	This is an estimate of the bandwidth of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
Bandwidth Min (Kb)	The minimum available bandwidth for the TCP connections on the links within this entity.
Est Round Trip Time (us)	This is an estimate of the round trip delay of the network pipe used for the B-D product computation, which lets us derive the TCP receive window to advertise.
PMTU Enable	The path MTU discovery.
PMTU ResetTimeout (sec)	The time interval for which the discovered path MTU is valid, before MSS reverts back to the negotiated TCP value.
CWM Enable	If true, congestion window monitoring is enabled.
CWM BurstSize (KB)	The maximum burst sent after a TCP sender idle period.
Max Jitter	The maximum delay variation (not due to congestion) that can be experienced by TCP connections on this interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

FCIP Tunnels

Field	Description
Interface	This identifies the interface on this FCIP device to which this link pertains.
Attached	The interface on which this FCIP link was initiated.
B Port Enable	If true, the B port mode is enabled on the local FCIP link.
B Port KeepAlive	If true, a message is sent in response to a (Fibre Channel) ELS Echo frame received from the peer. Some B Port implementations use ELS Echo request/response frames as Link Keep Alive.
Remote IP Address	The Internet address for the remote FCIP entity.
Remote TCP Port	The remote TCP port to which the local FCIP entity will connect if and when it initiates a TCP connection setup for this link.
Spc Frames Enable	If true, the TCP active opener initiates FCIP special frames and the TCP passive opener responds to the FCIP special frames. If it is set to false, the FCIP special frames are neither generated nor responded to.
Spc Frames RemoteWWN	The world wide name of the remote FC fabric entity. If this is a zero length string then this link would accept connections from any remote entity. If a WWN is specified then this link would accept connections from a remote entity with this WWN.
Spc Frames Remote Profile Id	The remote FCIP entity's identifier.

FCIP Tunnels (Advanced)

Field	Description
Interface	The interface on which this FCIP link was initiated.
Timestamp Enable	If true, the timestamp in FCIP header is to be checked.
Timestamp Tolerance	The accepted time difference between the local time and the timestamp value received in the FCIP header. By default this value will be EDTOV/2. EDTOV is the Error_Detect_Timeout Value used for Fibre Channel ports as the timeout value for detecting an error condition.
Number Connections	The maximum number of TCP connections allowed on this link.
Passive	If false, this link endpoint actively tries to connect to the peer. If true, the link endpoint waits for the peer to connect to it.
QoS Control	The value to be set for the ToS field in IP header for the TCP control connection.
QoS Data	The value to be set for the ToS field in IP header for the TCP data connection.
IP Compression	What algorithm is used, if any.
Write Accelerator	The write accelerator allows for enhancing SCSI write performance.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Tape Accelerator	If true, the tape accelerator (which allows for enhancing Tape write performance) is enabled.
Tape Accelerator Oper	Write acceleration is enabled for the FCIP link.
TapeRead Accelerator Oper	Enabled automatically when the tape accelerator oper is active.
FlowCtrlBufSize Tape (KB)	The size of the flow control buffer (64 K to 32 MB). If set to 0, flow control buffer size is calculated automatically by the switch.
IPSec	Indicates whether the IP security has been turned on or off on this link.
XRC Emulator	Check to enable XRC emulator. It is disabled by default.
XRC Emulator Oper	Indicates the operational status of XRC emulator.

FCIP Tunnels (FICON TA)

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
VSAN List Admin	The list of VSANs for which FICON tape acceleration is configured.
VSAN List Oper	The list of VSANs for which FICON tape acceleration is operationally on.

FCIP Tunnels Statistics

Field	Description
Interface	A unique value that identifies the interface on this FCIP device to which this link pertains.
Rx IPCompRatio	The IP compression ratio for received packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.
Tx IPCompRatio	The IP compression ratio for transmitted packets on the FCIP device. The value of this object will be presented as a floating point number with two digits after the decimal point.

FCIP XRC Statistics

Field	Description
ProfileId	Unique ID of the profile.
Interface	Name of the interface.
RRSAccelerated	The number of read record set IUs accelerated.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
RRSForwarded	Number of read record set IUs forwarded.
BusyStatus	Number of instances of busy status received from the control unit.
UnitCheckStatus	Number of instances of unit check status received from the control unit.
cfmFcipLinkExtXRCEStats SelReset	Number of selective resets processed.
BufferAllocErrors	Number of buffer allocation errors.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.
CID	The iSCSI connection ID for this connection.
State	The current state of this connection, from an iSCSI negotiation point of view. <ul style="list-style-type: none"> login— The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. full— A valid iSCSI login response with the final bit set has been sent or received. logout— A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. Note that the size of reported in bytes even though the negotiation is in 512 K blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Initiator Type	Indicates whether the node is a host that participates in iSCSI load-balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

Module Control

Field	Description
Module Id	ID of the module.
Admin Status	Enables or disables the iSCSI feature for the module.
OperStatus	Shows whether the iSCSI interface is enabled or disabled for the module.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of command PDUs transferred on this session.
PDU Response	The count of response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Initiator Specific Target

Field	Description
Name	A globally unique identifier for the node.
Port WWN(s) Primary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
Port WWN(s) Secondary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) iSCSI	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Primary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
LUN Map (Hex) FC Secondary	The Fibre Channel target's port addresses associated with this iSCSI initiator-specific target.
No AutoZone Creation	Indicates if a Fibre Channel zone is automatically created for this iSCSI initiator-target and the iSCSI initiator. If true the zone is not automatically created. If false (default) the zone is automatically created.
Trespass Mode	The trespass mode for this node. If true the Fibre Channel node instance presents all LUN I/O requests to the secondary port (fcSecondaryAddress) if the primary port (fcAddress) is down.
Revert to Primary Port	The revert to primary mode for this node. If true the Fibre Channel node instance presents all LUN I/O requests to the primary port (fcAddress) when the primary port comes back online.
Primary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.
Secondary PWWN VSAN	Indicates the VSAN into which the auto zone is placed for this initiator target. If this object is not set then the VSAN is determined by querying the name server.

iSCSI Initiator PWWN

Field	Description
Port WWN	The Fibre Channel address for this entry.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none"> normal—session is a normal iSCSI session discovery—session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

Additional References

For additional information related to implementing IP storage, see the following section:

- [Related Document, page 6-29](#)
- [Standards, page 6-29](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- RFCs, page 6-29
- MIBs, page 6-29

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	–

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	–

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs

Send documentation comments to dcnm-san-docfeedback@cisco.com



Configuring IPv4 for Gigabit Ethernet Interfaces

Cisco MDS 9000 Family switches support IP version 4 (IPv4) on Gigabit Ethernet interfaces. This chapter describes how to configure IPv4 addresses and other IPv4 features.

This chapter includes the following topics:

- [Information About IPv4, page 7-1](#)
- [Licensing Requirements for IPv4 for Gigabit Ethernet Interfaces, page 7-3](#)
- [Guidelines and Limitations, page 7-4](#)
- [Default Settings, page 7-4](#)
- [Configuring IPv4, page 7-4](#)
- [Verifying IPV4 Configuration, page 7-9](#)
- [Configuration Examples for IPV4, page 7-11](#)
- [Additional References, page 7-12](#)

Information About IPv4

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.



Note

The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.



Note

For information about configuring FCIP, see [Chapter 2, “Configuring FCIP.”](#) For information about configuring iSCSI, see [Chapter 4, “Configuring iSCSI.”](#)

Text Part Number:

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed through the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems do not require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.

**Note**

The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.

**Note**

To configure IPv6 on a Gigabit Ethernet interface, see the [“Configuring IPv6 Addressing and Enabling IPv6 Routing”](#) section on page 8-14.

**Tip**

Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port. They should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

This section includes the following topics:

- [Interface Descriptions, page 7-2](#)
- [Beacon Mode, page 7-2](#)
- [About VLANs for Gigabit Ethernet, page 7-2](#)
- [Interface Subnet Requirements, page 7-3](#)

Interface Descriptions

See the [Interfaces Configuration Guide, Cisco DCNM for SAN](#) *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for details on configuring the switch port description for any interface.

Beacon Mode

See the [Interfaces Configuration Guide, Cisco DCNM for SAN](#) *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for details on configuring the beacon mode for any interface.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name:

slot-number / port-number.VLAN-ID

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 7-1](#)).

Table 7-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN IDs cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A subinterface cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	

**Note**

The configuration requirements in [Table 7-1](#) also apply to Ethernet PortChannels.

Licensing Requirements for IPv4 for Gigabit Ethernet Interfaces

The following table shows the licensing requirements for this feature:

License	License Description
Enterprise package (ENTERPRISE_PKG)	It comprises IPsec and IKE for IPv4.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Guidelines and Limitations

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note

Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established** option is ignored when you apply IPv4-ACLs containing this option to Gigabit Ethernet interfaces.
 - If an IPv4-ACL rule applies to a pre-existing TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B and an IPv4-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.



Tip

If IPv4-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

Default Settings

Table 7-2 lists the default settings for IPv4 parameters.

Table 7-2 *Default IPv4 Parameters*

Parameters	Default
IPv4 MTU frame size	1500 bytes for all Ethernet ports
Autonegotiation	Enabled
Promiscuous mode	Disabled

Configuring IPv4

This section includes the following topics:

- [Configuring Gigabit Ethernet Interface, page 7-5](#)
- [Configuring Autonegotiation, page 7-5](#)
- [Configuring the MTU Frame Size, page 7-6](#)
- [Configuring Promiscuous Mode, page 7-7](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Configuring the VLAN Subinterface](#), page 7-7
- [Configuring Static IPv4 Routing](#), page 7-8
- [Applying IPv4-ACLs on Gigabit Ethernet Interfaces](#), page 7-8
- [Clearing ARP Cache](#), page 7-9

Configuring Gigabit Ethernet Interface

Detailed Steps

To configure the Gigabit Ethernet interface, follow these steps:

-
- Step 1** Expand **Switches > Interfaces > Ethernet > IPS**.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** Click the **IP Addresses** tab.
- Step 3** Click **Create Row**.
You see the Create Gigabit Ethernet Interface dialog box.
- Step 4** Select the switch on which you want to create the Gigabit Ethernet interface.
- Step 5** Enter the interface. For example, 2/2 for slot 2, port 2.
- Step 6** Enter the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0).
- Step 7** Click **Create** to save these changes or click **Close** to discard any unsaved changes.
-

Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

Detailed Steps

To configure autonegotiation, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# switchport auto-negotiate	Enables autonegotiation for this Gigabit Ethernet interface (default).
	switch(config-if)# no switchport auto-negotiate	Disables autonegotiation for this Gigabit Ethernet interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To configure autonegotiation, follow these steps:

-
- Step 1** Expand **Switches > Interfaces > Ethernet > IPS**.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** In the General tab, you can enable or disable the Auto Negotiate option for a specific switch.
- Step 3** Click **Apply Changes**.
-

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note The minimum MTU size is 576 bytes.



Tip MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

You do not need to explicitly issue the **shutdown** and **no shutdown** commands.

Detailed Steps

To configure the MTU frame size, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# switchport mtu 3000	Changes the MTU size to 3000 bytes. The default is 1500 bytes.

To configure the MTU frame size, follow these steps:

-
- Step 1** Expand **Switches > Interfaces > Ethernet > IPS**.
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** In the General tab, in the Mtu column, you can enter a new value to configure the MTU Frame Size for a specific switch. For example 3000 bytes. The default is 1500 bytes.
- Step 3** Click **Apply Changes**.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

Detailed Steps

To configure the promiscuous mode, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# switchport promiscuous-mode on	Enables promiscuous mode for this Gigabit Ethernet interface. The default is off .
	switch(config-if)# switchport promiscuous-mode off	Disables (default) promiscuous mode for this Gigabit Ethernet interface.
	switch(config-if)# no switchport promiscuous-mode	Disables (default) the promiscuous mode for this Gigabit Ethernet interface.

To configure the promiscuous mode, follow these steps:

-
- Step 1** [Expand Switches > Interfaces > Ethernet > IPS.](#)
You see the Gigabit Ethernet Configuration in the Information pane.
- Step 2** [In the General tab, you can enable or disable the Promiscuous Mode option for a specific switch.](#)
- Step 3** [Click Apply Changes.](#)
-

Configuring the VLAN Subinterface

Detailed Steps

To configure a VLAN subinterface (VLAN ID), follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2.100 switch(config-if)#	Specifies the subinterface on which 802.1Q is used (slot 2, port 2, VLAN ID 100). Note The subinterface number, 100 in this example, is the VLAN ID. The VLAN ID ranges from 1 to 4093.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 3	<code>switch(config-if)# ip address 10.1.1.101 255.255.255.0</code>	Enters the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	<code>switch(config-if)# no shutdown</code>	Enables the interface.

To configure a VLAN subinterface (VLAN ID) using Device Manager, follow these steps:

-
- Step 1 Select **Interface > Ethernet and iSCSI**.
 - Step 2 Click the **Sub Interfaces** tab.
 - Step 3 Select the Gigabit Ethernet subinterface on which 802.1Q should be used.
 - Step 4 Click the **Edit IP Address** button.
 - Step 5 Enter the IPv4 address and subnet mask for the Gigabit Ethernet interface.
 - Step 6 Click **Create** to save the changes or you may click **Close**.
-

Configuring Static IPv4 Routing

Detailed Steps

To configure static IPv4 routing (see Figure 7-1) through the Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	<code>switch# config terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1</code> <code>switch(config-if)#</code>	Enters the IP subnet (10.100.1.0 255.255.255.0) of the IP host and configures the next hop 10.1.1.1, which is the IPv4 address of the router connected to the Gigabit Ethernet interface.

Applying IPv4-ACLs on Gigabit Ethernet Interfaces

Detailed Steps

To apply an IPv4-ACL on a Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# interface gigabitethernet 3/1</code> <code>switch(config-if)#</code>	Configures a Gigabit Ethernet interface (3/1).

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command	Purpose
Step 3	<code>switch(config-if)# ip access-group SampleName</code>	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for both ingress and egress traffic (if the association does not exist already).
Step 4	<code>switch(config-if)# ip access-group SampleName1 in</code>	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for ingress traffic.
	<code>switch(config-if)# ip access-group SampleName2 out</code>	Applies the IPv4-ACL SampleName on Gigabit Ethernet 3/1 for egress traffic (if the association does not exist already).

Clearing ARP Cache

Detailed Steps

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache.

Use the **clear ips arp** command to clear the ARP cache. See Example 7-1 and Example 7-2.

Examples

Example 7-1 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 7-2 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```

Verifying IPV4 Configuration

To display IPv4 configuration information, perform one of the following tasks:

Command	Purpose
<code>ping 10.100.1.25</code>	Verifies gigabit ethernet connectivity.
<code>show ips ip route interface gig 8/1</code>	Displays the IP route table.
<code>show ips arp interface gigabitethernet 7/1</code>	Displays ARP caches.
<code>clear ips arp address 10.2.2.2 interface gigabitethernet 8/7</code>	Clears one ARP cache entry.
<code>clear ips arp interface gigabitethernet 8/7</code>	Clears all ARP cache entries.
<code>show ips stats ip interface gigabitethernet 4/1</code>	Displays IPv4 statistics.

Send documentation comments to dcnm-san-docfeedback@cisco.com

This section includes the following topics:

- Verifying Gigabit Ethernet Connectivity, page 7-10
- Displaying the IPv4 Route Table, page 7-10
- Displaying ARP Cache, page 7-11
- Displaying IPv4 Statistics, page 7-11

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.



Note

If the connection fails, verify the following, and ping the IP host again:

- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the up state.

Use the **ping** command to verify the Gigabit Ethernet connectivity (see Example 7-3). The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the “Using the ping and ping ipv6 Commands”).

Use the **show interface gigabitethernet** command to verify if the Gigabit Ethernet interface is up.

Example 7-3 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
```

Displaying the IPv4 Route Table

The **ip route interface** command takes the Gigabit Ethernet interface as a parameter and returns the route table for the interface. See Example 7-4.

Example 7-4 Displays the IP Route Table

```
switch# show 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.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying ARP Cache

You can display the ARP cache on Gigabit Ethernet interfaces.



Note

Use the physical interface, not the subinterface, for all ARP cache commands.

Use the **show ips arp interface gigabitethernet** command to display the ARP cache on the Gigabit Ethernet interfaces. This command takes the Ethernet interface as a parameter and returns the ARP cache for that interface. See Example 7-5.

Example 7-5 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
...
```

Displaying IPv4 Statistics

Use the **show ips stats ip interface gigabitethernet** to display and verify IP v4 statistics. This command takes the main Ethernet interface as a parameter and returns the IPv4 statistics for that interface. See Example 7-6.



Note

Use the physical interface, not the subinterface, to display IPv4 statistics.

Example 7-6 Displays IPv4 Statistics


```
switch# show ips stats ip interface gigabitethernet 4/1
Internet Protocol Statistics for port GigabitEthernet4/1
 168 total received, 168 good, 0 error
 0 reassembly required, 0 reassembled ok, 0 dropped after timeout
 371 packets sent, 0 outgoing dropped, 0 dropped no route
 0 fragments created, 0 cannot fragment
```

Configuration Examples for IPv4

Figure 7-1 shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 7-1 Gigabit Ethernet IPv4 Configuration Example



Note

The port on the Ethernet switch to which the MDS Gigabit Ethernet interface is connected should be configured as a host port (also known as access port) instead of a switch port. Spanning tree configuration for that port (on the Ethernet switch) should be disabled. This helps avoid the delay in the management port coming up due to delay from Ethernet spanning tree processing that the Ethernet switch would run if enabled. For Cisco Ethernet switches, use either the **switchport host** command in Cisco IOS or the **set port host** command in the Catalyst OS.

To configure the Gigabit Ethernet interface for the example in Figure 7-1, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.100 255.255.255.0	Enters the IPv4 address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

Additional References

For additional information related to implementing FCIPs, see the following section:

- [Related Document, page 7-13](#)
- [Standards, page 7-13](#)
- [RFCs, page 7-13](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 7-13](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-IETF-VRRP-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs

■ Additional References

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)



Configuring IPv6 for Gigabit Ethernet Interfaces

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS NX-OS. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.



Note

For Cisco NX-OS features that use IP addressing, refer to the chapters in this guide that describe those features for information on IPv6 addressing support.



Note

To configure IP version 4 (IPv4) on a Gigabit Ethernet interface, see [Chapter 7, “Configuring IPv4 for Gigabit Ethernet Interfaces.”](#)

This chapter includes the following topics:

- [Information About IPV6, page 8-1](#)
- [Guidelines and Limitations, page 8-12](#)
- [Default Settings, page 8-13](#)
- [Configuring Basic Connectivity for IPv6, page 8-13](#)
- [Verifying IPV6 Configuration, page 8-18](#)
- [Configuration Examples for IPV6, page 8-21](#)
- [Additional References, page 8-23](#)

Information About IPV6

IP version 6 (IPv6) provides extended addressing capability beyond those provided in IP version 4 (IPv4) in Cisco MDS NX-OS by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits. The architecture of IPv6 has been designed to allow existing IPv4 users to transition easily to IPv6 while providing services such as end-to-end security, quality of service (QoS), and globally unique addresses.

IPv6 provides the following enhancements over IPv4:

- Allows networks to scale and provide global reachability.
- Reduces the need for private address and network address translation (NAT).

Text Part Number:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Provides simpler autoconfiguration of addresses.

This section describes the IPv6 features supported by Cisco MDS NX-OS and includes the following topics:

- [Extended IPv6 Address Space for Unique Addresses, page 8-2](#)
- [IPv6 Address Formats, page 8-2](#)
- [IPv6 Address Prefix Format, page 8-3](#)
- [IPv6 Address Type: Unicast, page 8-3](#)
- [IPv6 Address Type: Multicast, page 8-5](#)
- [ICMP for IPv6, page 8-6](#)
- [Path MTU Discovery for IPv6, page 8-7](#)
- [IPv6 Neighbor Discovery, page 8-7](#)
- [Router Discovery, page 8-9](#)
- [IPv6 Stateless Autoconfiguration, page 8-9](#)
- [Dual IPv4 and IPv6 Protocol Stacks, page 8-10](#)
- [IPv6 Addressing and Enabling IPv6 Routing, page 8-11](#)
- [Transitioning from IPv4 to IPv6, page 8-12](#)

Extended IPv6 Address Space for Unique Addresses

IPv6 extends the address space by quadrupling the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides many more globally unique IP addresses. By being globally unique, IPv6 addresses enable global reachability and end-to-end security for networked devices, functionality that is crucial to the applications and services that are driving the demand for more addresses.

IPv6 Address Formats

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format `x:x:x:x:x:x:x`. The following are examples of IPv6 addresses:

```
2001:0DB8:7654:3210:FEDC:BA98:7654:3210
```

```
2001:0DB8:0:0:8:800:200C:417A
```

It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (:) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). [Table 8-1](#) lists compressed IPv6 address formats.



Note

Two colons (:) can be used only once in an IPv6 address to represent the longest successive hexadecimal fields of zeros.



Note

The hexadecimal letters in IPv6 addresses are not case-sensitive.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 8-1 Compressed IPv6 Address Formats

IPv6 Address Type	Uncompressed Format	Compressed Format
Unicast	2001:0DB8:800:200C:0:0:0:417A	2001:0DB8:800:200C::417A
Multicast	FF01:0:0:0:0:0:101	FF01::101

IPv6 Address Prefix Format

An IPv6 address prefix, in the format *ipv6-prefix/prefix-length*, can be used to represent bit-wise contiguous blocks of the entire address space. The *ipv6-prefix* is specified in hexadecimal using 16-bit values between the colons. The *prefix-length* is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Address Type: Unicast


An IPv6 unicast address is an identifier for a single interface on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address. The Cisco MDS NX-OS supports the following IPv6 unicast address types:

- Global addresses
- Link-local addresses

Global Addresses

Global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Figure 8-1 shows the structure of a global address.

Figure 8-1 Global Address Format



Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.

The aggregatable global address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA). The IETF decided to remove the TLA and NLA fields from the RFCs because these fields are policy-based. Some existing IPv6 networks deployed before the change might still be using networks based on the older architecture.


Send documentation comments to dcnm-san-docfeedback@cisco.com

A 16-bit subnet field called the subnet ID could be used by individual organizations to create their own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

An interface ID is used to identify interfaces on a link. The interface ID must be unique to the link. They may also be unique over a broader scope. In many cases, an interface ID will be the same as, or based on, the link-layer address of an interface, which results in a globally unique interface ID. Interface IDs used in aggregatable global unicast and other IPv6 address types must be 64 bits long and constructed in the modified EUI-64 format.

Cisco MDS NX-OS supports IEEE 802 interface types (for example, Gigabit Ethernet interfaces). The first three octets (24 bits) are taken from the Organizationally Unique Identifier (OUI) of the 48-bit link-layer address (MAC address) of the interface, the fourth and fifth octets (16 bits) are a fixed hexadecimal value of FFFE, and the last three octets (24 bits) are taken from the last three octets of the MAC address. The construction of the interface ID is completed by setting the Universal/Local (U/L) bit—the seventh bit of the first octet—to a value of 0 or 1. A value of 0 indicates a locally administered identifier; a value of 1 indicates a globally unique IPv6 interface identifier (see Figure 8-2).

Figure 8-2 Interface Identifier Format




Link-Local Address

A link-local address is an IPv6 unicast address that is automatically configured on an interface using the link-local prefix FE80::/10 and the interface identifier in the modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate. Figure 8-3 shows the structure of a link-local address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 8-3 Link-Local Address Format

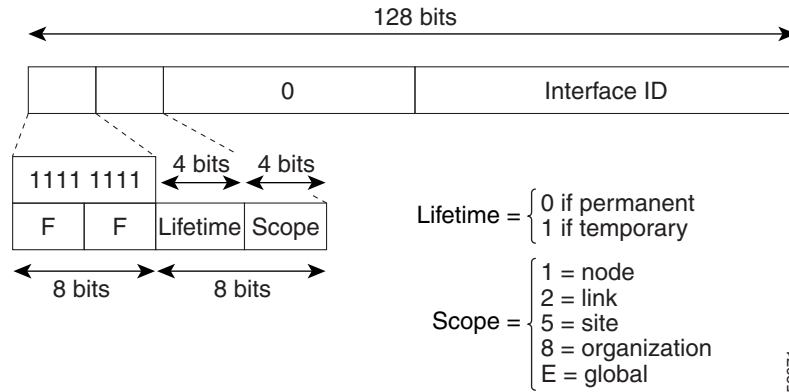


52669

IPv6 Address Type: Multicast

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure 8-4 shows the format of the IPv6 multicast address.

Figure 8-4 IPv6 Multicast Address Format



52671

IPv6 hosts are required to join (receive packets destined for) the following multicast groups:


- All-node multicast group FF02::1.
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 concatenated with the low-order 24 bit of the unicast address.

The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6

Send documentation comments to dcnm-san-docfeedback@cisco.com

unicast address. (See Figure 8-5) For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C. Solicited-node addresses are used in neighbor solicitation messages.

Figure 8-5 IPv6 Solicited-Node Multicast Address Format



Note

There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.


ICMP for IPv6

Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4—ICMP generates error messages such as ICMP destination unreachable messages, and informational messages such as ICMP echo request and reply messages. Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery process, path MTU discovery, and the Multicast Listener Discovery (MLD) protocol for IPv6. MLD is based on version 2 of the Internet Group Management Protocol (IGMP) for IPv4.

A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet. ICMP packets in IPv6 resemble a transport-layer packet in the sense that the ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet. Within IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is derived (computed by the sender and checked by the receiver) from the fields in the IPv6 ICMP packet and the IPv6 pseudoheader. The ICMPv6 Data field contains error or diagnostic information relevant to IP packet processing. Figure 8-6 shows the IPv6 ICMP packet header format.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 8-6 IPv6 ICMP Packet Header Format



Path MTU Discovery for IPv6

As in IPv4, path MTU discovery in IPv6 allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, however, fragmentation is handled by the source of a packet when the path MTU of one link along a given data path is not large enough to accommodate the size of the packets. Having IPv6 hosts handle packet fragmentation saves IPv6 router processing resources and helps IPv6 networks run more efficiently.



Note

In IPv4, the minimum link MTU is 68 octets, which means that the MTU size of every link along a given data path must support an MTU size of at least 68 octets.

In IPv6, the minimum link MTU is 1280 octets. We recommend using MTU value of 1500 octets for IPv6 links.

IPv6 Neighbor Discovery

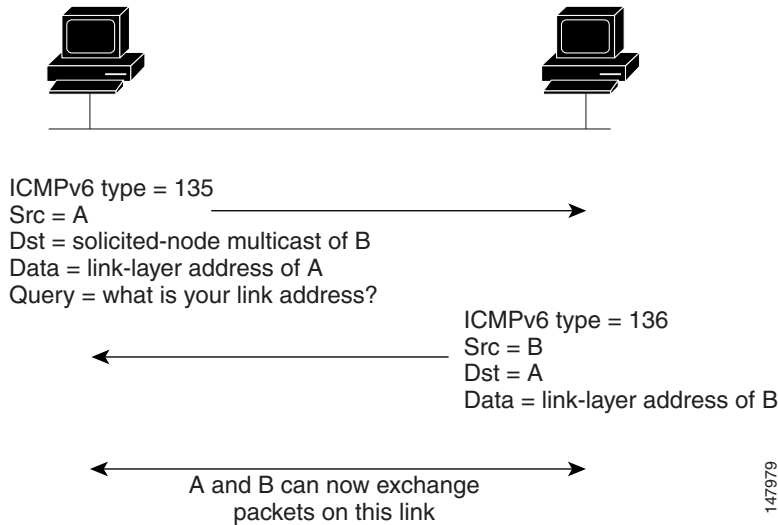
The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers.

IPv6 Neighbor Solicitation and Advertisement Messages

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. (See Figure 8-7.) When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 8-7 IPv6 Neighbor Discovery—Neighbor Solicitation Message



After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verify the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-node multicast address.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor, and is used for all paths between hosts and neighboring nodes (hosts or routers). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when the neighbor returns a positive acknowledgment indicating that it has received and processed packets previously sent to it. A positive acknowledgment could be from an upper-layer protocol such as TCP indicating that a connection is making forward progress (reaching its destination) or the receipt of a neighbor advertisement message in response to a neighbor solicitation message. If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Therefore, forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop router is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working. The return of a solicited neighbor advertisement message from the neighbor is a positive

Send documentation comments to dcnm-san-docfeedback@cisco.com

acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.

**Note**

A neighbor advertisement message that has the solicited flag set to a value of 0 must not be considered as a positive acknowledgment that the forward path is still working.

Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). Specifically, a node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

Every IPv6 unicast address (global or link-local) must be checked for uniqueness on the link; however, until the uniqueness of the link-local address is verified, duplicate address detection is not performed on any other IPv6 addresses associated with the link-local address.

Router Discovery

Router discovery performs both router solicitation and router advertisement. Router solicitations are sent by hosts to all-routers multicast addresses. Router advertisements are sent by routers in response to solicitations or unsolicited and contain default router information as well as additional parameters such as the MTU and hop limit.


IPv6 Stateless Autoconfiguration

All interfaces on IPv6 nodes must have a link-local address, which is automatically configured from the identifier for an interface and the link-local prefix FE80::/10. A link-local address enables a node to communicate with other nodes on the link and can be used to further configure the node.

Nodes can connect to a network and automatically generate site-local and global IPv6 address without the need for manual configuration or help of a server, such as a DHCP server. With IPv6, a router on the link advertises in router advertisement (RA) messages any site-local and global prefixes, and its willingness to function as a default router for the link. RA messages are sent periodically and in response to router solicitation messages, which are sent by hosts at system startup. (See Figure 8-8.)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 8-8 IPv6 Stateless Autoconfiguration




A node on the link can automatically configure site-local and global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure their uniqueness on the link. If the prefixes advertised in the RA messages are globally unique, then the IPv6 addresses configured by the node are also guaranteed to be globally unique. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message.

Dual IPv4 and IPv6 Protocol Stacks

The dual IPv4 and IPv6 protocol stack technique is one technique for a transition to IPv6. It enables gradual, one-by-one upgrades to applications running on nodes. Applications running on nodes are upgraded to make use of the IPv6 protocol stack. Applications that are not upgraded—they support only the IPv4 protocol stack—can coexist with upgraded applications on the same node. New and upgraded applications simply make use of both the IPv4 and IPv6 protocol stacks. (See Figure 8-9.)

Figure 8-9 Dual IPv4 and IPv6 Protocol Stack Technique




A new API has been defined to support both IPv4 and IPv6 addresses and DNS requests. An application can be upgraded to the new API and still use only the IPv4 protocol stack. The Cisco MDS NX-OS supports the dual IPv4 and IPv6 protocol stack technique. When an interface is configured with both an IPv4 and an IPv6 address, the interface will accept and process both IPv4 and IPv6 traffic.

Send documentation comments to dcnm-san-docfeedback@cisco.com

In Figure 8-10, an application that supports dual IPv4 and IPv6 protocol stacks requests all available addresses for the destination host name `www.a.com` from a DNS server. The DNS server replies with all available addresses (both IPv4 and IPv6 addresses) for `www.a.com`. The application chooses an address—in most cases, IPv6 addresses are the default choice—and connects the source node to the destination using the IPv6 protocol stack.

Figure 8-10 Dual IPv4 and IPv6 Protocol Stack Applications



IPv6 Addressing and Enabling IPv6 Routing

IPv6 addresses are represented as a series of 16-bit hexadecimal fields separated by colons (:) in the format `x:x:x:x:x:x:x`. It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses easier to use, two colons (:) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). By default, IPv6 addresses are not configured, and IPv6 processing is disabled. You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group `FF02:0:0:0:0:1:FF00::/104` for each unicast address assigned to the interface
- All-node link-local multicast group `FF02::1`

This task explains how to assign IPv6 addresses to individual router interfaces and enable the processing of IPv6 traffic. By default, IPv6 addresses are not configured and IPv6 processing is disabled.

You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN

Send documentation comments to dcnm-san-docfeedback@cisco.com



Note

The IPv6 address *ipv6-address* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The IPv6 prefix *ipv6-prefix* argument in the **ipv6 address** command must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons.

The IPv6 prefix length *prefix-length* argument in the **ipv6 address** command is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.

Configuring a global IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Additionally, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:1:FF00::/104 for each unicast address assigned to the interface
- All-node link-local multicast group FF02::1



Note

The solicited-node multicast address is used in the neighbor discovery process.



Note

The maximum number of IPv6 addresses (static and autoconfigured) allowed on an interface is eight, except on the management (mgmt 0) interface where only one static IPv6 address can be configured.

Transitioning from IPv4 to IPv6

Cisco MDS NX-OS does not support any transitioning mechanisms from IPv4 to IPv6. However, you can use the transitioning schemes in the Cisco router products for this purpose. For information on configuring Cisco routers to transition your network, refer to the “[Implementing Tunneling for IPv6](#)” chapter in the *Cisco IOS IPv6 Configuration Guide*.

Guidelines and Limitations



Tip

If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See the *Security Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS Security Configuration Guide for information on configuring IPv6-ACLs.

Follow these guidelines when configuring IPv6-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



Note

Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- Apply IPv6-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established** option is ignored when you apply IPv6-ACLs containing this option to Gigabit Ethernet interfaces.
 - If an IPv6-ACL rule applies to a preexisting TCP connection, that rule is ignored. For example, if there is an existing TCP connection between A and B and an IPv6-ACL that specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

See the *Security Configuration Guide, Cisco DCNM for SAN* Cisco MDS 9000 Family NX-OS Security Configuration Guide for information on applying IPv6-ACLs to an interface.

Default Settings

Table 8-2 lists the default settings for IPv6 parameters.

Table 8-2 **Default IPv6 Parameters**

Parameters	Default
IPv6 processing	Disabled
Duplicate address detection attempts	0 (neighbor discovery disabled)
Reachability time	1000 milliseconds
Retransmission time	30000 milliseconds
IPv6-ACLs	None

Configuring Basic Connectivity for IPv6

This section includes the following topics:

- [Configuring IPv6 Addressing and Enabling IPv6 Routing, page 8-14](#)
- [Configuring IPv4 and IPv6 Protocol Addresses, page 8-15](#)
- [Clearing IPv6 Neighbor Discovery Cache, page 8-16](#)
- [Configuring Neighbor Discovery Parameters, page 8-16](#)
- [Configuring a IPv6 Static Route, page 8-17](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring IPv6 Addressing and Enabling IPv6 Routing

Detailed Steps

To configure an IPv6 address on an interface and enable IPv6 routing, follow these steps:

	Command or Action	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 1/1 switch(config-if)#	Specifies a Gigabit Ethernet interface and enters interface configuration submode.
	switch(config)# interface mgmt 0 switch(config-if)#	Specifies the management interface and enters interface configuration submode.
	switch(config)# interface gigabitethernet 2/2.100 switch(config-if)#	Specifies a Gigabit Ethernet subinterface (VLAN ID) and enters interface configuration submode.
	switch(config)# interface vsan 10 switch(config-if)#	Specifies a VSAN interface and enters interface configuration submode.
Step 3	switch(config-if)# ipv6 address 2001:0DB8:800:200C::417A/64	Assigns a unicast IPv6 address to the interface, automatically configures an IPv6 link-local address on the interface, and enables IPv6 processing on the interface.
	switch(config-if)# ipv6 address autoconfig	Enables autoconfiguration of the IPv6 link-local and unicast addresses on the interface, and enables IPv6 processing on the interface.
	switch(config-if)# ipv6 enable	Automatically configures an IPv6 link-local address on the interface while also enabling the interface for IPv6 processing. The link-local address can be used only to communicate with nodes on the same link.
Step 4	switch(config-if)# no shutdown	Enables the interface.
Step 5	switch(config-if)# exit switch(config)	Exits interface configuration submode and returns to configuration mode.
Step 6	switch(config)# ipv6 routing	Enables the processing of IPv6 unicast datagrams.

To configure an IPv6 address on an interface using Device Manager, follow these steps:

-
- Step 1** Choose **Interfaces > Gigabit Ethernet and iSCSI**.
You see the Gigabit Ethernet Configuration dialog box.
 - Step 2** Click the IP Address that you want to configure and click **Edit IP Address**.
You see the IP Address dialog box.
 - Step 3** Click **Create** and set the IP Address/Mask field, using the IPv6 format (for example, 2001:0DB8:800:200C::417A/64).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click **Create** to save these changes or click **Close** to discard any unsaved changes.

To enable IPv6 routing using Device Manager, follow these steps:

Step 1 Choose **IP > Routing**. You see the IP Routing Configuration dialog box.

Step 2 Check the **Routing Enabled** check box.

Step 3 Click **Apply** to save these changes or click **Close** to discard any unsaved changes.

Configuring IPv4 and IPv6 Protocol Addresses

When an interface in a Cisco networking device is configured with both an IPv4 and an IPv6 address, the interface can send and receive data on both IPv4 and IPv6 networks.

Detailed Steps

To configure an interface in a Cisco networking device to support both the IPv4 and IPv6 protocol stacks, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 1/1 switch(config-if)#	Specifies the interface, and enters interface configuration submenu.
Step 3	switch(config-if)# ip address 192.168.99.1 255.255.255.0	Specifies a primary or secondary IPv4 address for an interface.
Step 4	switch(config-if)# ipv6 address 2001:0DB8:c18:1::3/64	Specifies the IPv6 network assigned to the interface and enables IPv6 processing on the interface. Note See the “ Configuring IPv6 Addressing and Enabling IPv6 Routing ” section for more information on configuring IPv6 addresses.
Step 5	switch(config-if)# no shutdown	Enables the interface.
Step 6	switch(config-if)# exit switch(config)	Exits interface configuration submenu, and returns to configuration mode.
Step 7	switch(config)# ipv6 routing	Enables the processing of IPv6 unicast datagrams.

To configure an interface in a Cisco networking device to support both the IPv4 and IPv6 protocol stacks using Device Manager, follow these steps:

Step 1 Choose **Interfaces > Gigabit Ethernet and iSCSI**.

You see the Gigabit Ethernet Configuration dialog box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Click the IP Address field that you want to configure and click **Edit IP Address**.
You see the IP Address dialog box.
- Step 3** Click **Create** and set the IP Address/Mask field, using the IPv4 or IPv6 format.
- Step 4** Click **Create** to save these changes or click **Close** to discard any unsaved changes.

Clearing IPv6 Neighbor Discovery Cache

Detailed Steps

You can clear the IPv6 neighbor discovery cache using the **clear ipv6 neighbor** command in EXEC mode.

```
switch# clear ipv6 neighbor
```

Configuring Neighbor Discovery Parameters

You can configure the following neighbor discovery parameters:

- Duplicate address detection attempts
- Reachability time
- Retransmission timer



Note

We recommend that you use the factory-defined defaults for these parameters.

This section includes the following topics:

- Duplicate Address Detection Attempts, page 8-16
- Reachability Time, page 8-17
- Retransmission Time, page 8-17

Duplicate Address Detection Attempts

Detailed Steps

To configure the number of duplicate address detection attempts, follow these steps:

	Command or Action	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 3/1 switch(config-if)#	Specifies an interface and enters interface configuration submenu.

Send documentation comments to dcnm-san-docfeedback@cisco.com

	Command or Action	Purpose
Step 3	switch(config-if)# ipv6 nd dad attempts 3	Sets the duplicate address detection attempts count to 100. The range is 0 to 15.
Step 4	switch(config-if)# no ipv6 nd dad attempts	Reverts to the default value (0). Note When the attempt count is set to 0, neighbor discovery is disabled.

Reachability Time

Detailed Steps

To configure the reachability time, follow these steps:

	Command or Action	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 3/1 switch(config-if)#	Specifies an interface and enters interface configuration submenu.
Step 3	switch(config-if)# ipv6 nd reachability-time 10000	Sets the reachability time to 10000 milliseconds. The range is 1000 to 3600000 milliseconds.
Step 4	switch(config-if)# no ipv6 nd reachability-time	Reverts to the default value (30000 milliseconds).

Retransmission Time

Detailed Steps

To configure the retransmission time, follow these steps:

	Command or Action	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 3/1 switch(config-if)#	Specifies an interface and enters interface configuration submenu.
Step 3	switch(config-if)# ipv6 nd retransmission-timer 20000	Sets the retransmission time to 20000 milliseconds. The range is 1000 to 3600000 milliseconds.
Step 4	switch(config-if)# no ipv6 nd retransmission-timer	Reverts to the default value (1000 milliseconds).

Configuring a IPv6 Static Route

You must manually configure IPv6 static routes and define an explicit path between two networking devices. IPv6 static routes are not automatically updated and must be manually reconfigured if the network topology changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To configure a IPv6 static route, follow these steps:

	Command or Action	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ipv6 route ::/0 gigabitethernet 3/1	Configures a static default IPv6 route on a Gigabit Ethernet interface.
Step 3	switch(config)# ipv6 route 2001:0DB8::/32 gigabitethernet 3/2	Configures a fully specified IPv6 static route on a Gigabit Ethernet interface.

To configure a IPv6 static route using Device Manager, follow these steps:

-
- Step 1** Choose **IP > Routing**.
You see the IP Routing Configuration dialog box.
- Step 2** Click **Create**.
You see the Create IP Route dialog box.
- Step 3** Set the **Dest** field to the IPv6 destination address.
- Step 4** Set the **Mask** field to the IPv6 subnet mask.
- Step 5** Set the **Gateway** field to the IPv6 default gateway.
- Step 6** (Optional) Set the **Metric** field to the desired route metric.
- Step 7** Select the interface from the **Interface** drop-down menu.
- Step 8** Click **Create** to save these changes or click **Close** to discard any unsaved changes.
-

Verifying IPV6 Configuration

To display IPv6 configuration information, perform one of the following tasks:

Command	Purpose
switch# show ipv6 interface mgmt 0	Verifies that IPv6 addresses are configured correctly for the Gigabit Ethernet 6/1 interface.
switch# show ipv6 neighbours	Displays IPv6 neighbor discovery cache information for all interfaces.
switch# show ipv6 traffic	Displays IPv6 and ICMP statistics.
switch# show ipv6 interface mgmt 0	Displays the configuration of the neighbor discovery parameters.
switch# show ipv6 route	Displays the IPv6 route table for the switch.
show ips ipv6 neighbours interface gigabitethernet 6/1	Displays information about IPv6 neighbors for an interface.
show ips ipv6 prefix-list interface gigabitethernet 6/1	Displays information about IPv6 prefixes for an interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Command	Purpose
<code>show ips ipv6 route interface gigabitethernet 6/1</code>	Displays information about the IPv6 routes for an interface.
<code>show ips ipv6 routers interface gigabitethernet 6/1</code>	Displays information about IPv6 routers for an interface.
<code>show ips ipv6 traffic interface gigabitethernet 6/1</code>	Displays information about IPv6 traffic statistics for an interface.

This section contains the following topics:

- Verifying Neighbor Discovery Parameter Configuration, page 8-19
- Verifying IPv6 Static Route Configuration and Operation, page 8-19
- Displaying IPv6, page 8-20

Verifying Neighbor Discovery Parameter Configuration

The `show ipv6 interface` command displays the configuration of the neighbor discovery parameters.

```
switch# show ipv6 interface mgmt 0
mgmt0 is up
  IPv6 is enabled
  Global address(es):
    2003::1/64
  Link-local address(es):
    fe80::205:30ff:fe00:533e
  ND DAD is enabled, number of DAD attempts: 5
  ND reachable time is 50000 milliseconds
  ND retransmission time is 3000 milliseconds
  Stateless autoconfig for addresses disabled
```

Verifying IPv6 Static Route Configuration and Operation

The `show ipv6 route` command displays the IPv6 route table for the switch.

```
switch# show ipv6 route

IPv6 Routing Table
Codes: C - Connected, L - Local, S - Static G - Gateway
G    ::/0
      via fe80::211:5dff:fe53:500a, GigabitEthernet6/1, distance 2
G    ::/0
      via fe80::2d0:3ff:fe61:4800, mgmt0, distance 2
C    2000::/64
      via ::, mgmt0
C    2172:22::/64
      via ::, mgmt0, distance 2
C    3000:3::/64
      via fe80::205:30ff:fe01:7ed6, GigabitEthernet4/1
C    3000:4::/64
      via fe80::205:30ff:fe01:7ed6, GigabitEthernet4/1.250
C    3000:5::/64
      via fe80::213:1aff:fee5:e69b, GigabitEthernet5/4
C    3000:6::/64
      via fe80::213:1aff:fee5:e69b, GigabitEthernet5/4.250
```

Send documentation comments to dcnm-san-docfeedback@cisco.com

```

C    3000:7::/64
    via fe80::205:30ff:fe01:7ed7, GigabitEthernet4/2
C    3000:8::/64
    via fe80::205:30ff:fe01:7ed7, GigabitEthernet4/2.250
C    3000:9::/64
    via fe80::213:1aff:fee5:e69e, port-channel 3
C    3000:10::/64
    via fe80::213:1aff:fee5:e69e, port-channel 3.250
C    5000:1::/64
    via fe80::205:30ff:fe01:3917, GigabitEthernet6/2
C    5000:1::/64
    via fe80::205:30ff:fe01:3918, port-channel 4
C    6000:1:1:1::/64
    via fe80::205:30ff:fe01:3916, GigabitEthernet6/1
C    7000:1::/64
    via fe80::205:30ff:fe01:3917, GigabitEthernet6/2.250
C    7000:1::/64
    via fe80::205:30ff:fe01:3918, port-channel 4.250
C    7000:1:1:1::/64
    via fe80::205:30ff:fe01:3917, GigabitEthernet6/2, distance 2
L    fe80::/10
    via ::
L    ff00::/8
    via ::

```

Displaying IPv6

Use the **show ips ipv6 neighbours interface** command for information about IPv6 neighbors for an interface.

```

switch# show ips ipv6 neighbours interface gigabitethernet 6/1
IPv6 Address                               Age (min)  Link-layer Addr  State  Interface
fe80::211:5dff:fe53:500a                   0          0011.5d53.500a   S      Gigabi tEthernet6/1

```

Use the **show ips ipv6 prefix-list interface** command for information about IPv6 prefixes for an interface.

```

switch# show ips ipv6 prefix-list interface gigabitethernet 6/1
Prefix                                     Prefix-len  Addr
Valid Preferred
6000:1:1:1::                               64         ::
      2592000      604800

```

Use the **show ips ipv6 route interface** command for information about the IPv6 routes for an interface.

```

switch# show ips ipv6 route interface gigabitethernet 6/1
IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, G - Gateway, M - Multicast
C 6000:1:1:1::/64 is directly connected, GigabitEthernet6/1
C 6000:1:1:1::/64 is directly connected, GigabitEthernet6/1
C fe80::/64 is directly connected, GigabitEthernet6/1
M ff02::/32 is multicast, GigabitEthernet6/1
G ::/0 via fe80::211:5dff:fe53:500a, GigabitEthernet6/1

```

Use the **show ips ipv6 routers interface** command for information about IPv6 routers for an interface.

```

switch# show ips ipv6 routers interface gigabitethernet 6/1
Addr                               Lifetime  Expire
fe80::211:5dff:fe53:500a          1800     1781

```

Send documentation comments to dcnm-san-docfeedback@cisco.com

Use the **show ips ipv6 traffic interface** command for information about IPv6 traffic statistics for an interface.

```
switch# show ips ipv6 traffic interface gigabitethernet 6/1
IPv6 statistics:
  Rcvd: 5094 total
        0 bad header, 0 unknown option, 0 unknown protocol
        0 fragments, 0 total reassembled
        0 reassembly timeouts, 0 reassembly failures
  Sent: 13625 generated
        0 fragmented into 0 fragments, 0 failed
        2 no route
ICMP statistics:
  Rcvd: 1264 input, 0 checksum errors, 0 too short
        0 unknown info type, 0 unknown error type
  unreachable: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
  parameter: 0 error, 0 header, 0 option
  0 hopcount expired, 0 reassembly timeout, 0 too big
  0 echo request, 0 echo reply
  734 group query, 0 group report, 0 group reduce
  0 router solicit, 528 router advert, 0 redirects
  0 neighbor solicit, 2 neighbor advert
  Sent: 6045 output, 0 rate-limited
  unreachable: 0 routing, 0 admin, 0 neighbor, 1160 address, 0 port
  parameter: 0 error, 0 header, 0 option
  0 hopcount expired, 0 reassembly timeout, 0 too big
  0 echo request, 0 echo reply
  0 group query, 1466 group report, 0 group reduce
  1 router solicit, 0 router advert, 0 redirects
  3412 neighbor solicit, 6 neighbor advert
```

Configuration Examples for IPv6

You can display information to verify the configuration and operation of basic IPv6 connectivity.

This section provides the following **show ipv6** command output examples:

- Example Output for the show ipv6 interface Command
- Example Output for the show ipv6 neighbours Command
- Example Output for the show ipv6 traffic Command

Example Output for the show ipv6 interface Command

In the following example, the **show ipv6 interface** command is used to verify that IPv6 addresses are configured correctly for the Gigabit Ethernet 6/1 interface.

```
switch# show ipv6 interface mgmt 0
mgmt0 is up
  IPv6 is enabled
  Global address(es):
    2172:22::180/64
  Link-local address(es):
    fe80::b8db:adff:feba:d074
  ND DAD is disabled
  ND reachable time is 30000 milliseconds
  ND retransmission time is 1000 milliseconds
  Stateless autoconfig for addresses disabled
  MTU is 1500 bytes
```

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Example Output for the show ipv6 neighbours Command

In the following example, the **show ipv6 neighbours** command displays IPv6 neighbor discovery cache information for all interfaces.

```
switch# show ipv6 neighbours
R - Reachable, I - Incomplete, S - Stale, F - Failed, P - Probe, D - Delay
IPv6 Address                               Age  State Link-layer Addr  Interface
fe80::211:5dff:fe53:500a                   0    S    0011.5d53.500a    GigE6/1
fe80::211:5dff:fe53:500a                   0    S    0011.5d53.500a    GigE6/2
5000:1::250                                 0    S    0011.5d53.500a    po 4
fe80::211:5dff:fe53:500a                   0    S    0011.5d53.500a    po 4
fe80::211:5dff:fe53:500a                   0    S    0011.5d53.500a    po 4
fe80::2d0:3ff:fe61:4800                     184  S    00d0.0361.4800    mgmt0
```

In the following example, the **show ipv6 neighbours interface** command displays IPv6 neighbor discovery cache information for the Gigabit Ethernet 6/1 interface.

```
switch# show ipv6 neighbours interface gigabitethernet 6/1
R - Reachable, I - Incomplete, S - Stale, F - Failed, P - Probe, D - Delay
IPv6 Address                               Age  State Link-layer Addr  Interface
fe80::211:5dff:fe53:500a                   0    S    0011.5d53.500a    GigE6/1
```

Example Output for the show ipv6 traffic Command

The **show ipv6 traffic** command displays IPv6 and ICMP statistics.

```
switch# show ipv6 traffic
IPv6 Statistics:
  Rcvd:  100 total, 0 local destination
         0 errors, 0 truncated, 0 too big
         0 unknown protocol, 0 dropped
         0 fragments, 0 reassembled
         0 couldn't reassemble, 0 reassembly timeouts
  Sent:  0 generated, 0 forwarded 0 dropped
         0 fragmented, 0 fragments created, 0 couldn't fragment

ICMPv6 Statistics:
  Rcvd:  100 total, 0 errors, 0 unreachable, 0 time exceeded
         0 too big, 0 param probs, 0 admin prohibits
         0 echos, 0 echo reply, 0 redirects
         0 group query, 0 group report, 0 group reduce
         0 router solicit, 69 router advert
         0 neighbor solicit, 31 neighbor advert
  Sent:  55 total, 0 errors, 0 unreachable, 0 time exceeded
         0 too big, 0 param probs, 0 admin prohibits
         0 echos, 0 echo reply, 0 redirects
         0 group query, 20 group report, 2 group reduce
         0 router solicit, 0 router advert
         0 neighbor solicit, 33 neighbor advert
```


Send documentation comments to dcnm-san-docfeedback@cisco.com

Additional References

For additional information related to implementing FCIPs, see the following section:

- [Related Document, page 8-23](#)
- [Standards, page 8-23](#)
- [RFCs, page 8-23](#)
- [MIBs, page 8-23](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference, Release 5.0(1a)</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-IETF-VRRP-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs

■ Additional References

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)



Symbols

* (asterisk)

iSCSI node [4-99](#)

A

AAA authentication

 configuring [4-50, 4-51](#)

access control

 enforcing iSCSI

 enforcing access control [4-13](#)

 iSCSI [4-11, 4-12](#)

access control zoning based access control iSCSI

 zoning based access control [4-13](#)

ACL based access control

 configuring for iSCSI [4-12](#)

ACLs

 configuring for iSCSI [4-12](#)

advertised interfaces [4-39](#)

advertisement packets

 setting time intervals [5-20](#)

ARP

 clearing entries [5-27](#)

 displaying entries [5-27](#)

ARP caches

 clearing [7-9](#)

 displaying [7-11](#)

authentication

 CHAP option [4-72](#)

 configuring local with Device Manager [4-52](#)

 iSCSI setup [4-71](#)

 local [4-52](#)

Text Part Number:

MD5 [5-22](#)

 mechanism [4-50](#)

 mutual CHAP mutual CHAP authentication [4-53](#)

 restricting iSLB initiator initiator authentication

 restricting iSLB

 restricting iSLB initiators [4-64](#)

 simple text [5-22](#)

 See also MD5 authentication

 See also simple text authentication

autogenerated iSCSI target iSCSI

 autogenerated target [4-13](#)

auto-negotiation

 configuring Gigabit Ethernet interfaces [6-6, 7-5](#)

B

B ports

 configuring [2-34](#)

 interoperability mode [2-9](#)

 SAN extenders [2-10](#)

bridge ports. See B ports

buffer sizes

 configuring in FCIP profiles [2-30](#)

C

CFS

 iSLB config distribution [4-21](#)

CHAP authentication [4-13, 4-19, 4-72](#)

 configuring for iSCSI [4-72](#)

CHAP challenge [4-53](#)

CHAP response [4-53](#)

CHAP user name [4-52](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Cisco Discovery Protocol [6-10](#)

Cisco Transport Controller. See CTC

cloud discovery. See iSNS cloud discovery

congestion window monitoring. See CWM

core dumps

IPS modules [6-5](#)

CTC

description [2-24](#)

launching [2-24](#)

Cut-through routing mode [4-15](#)

cut-thru routing mode [4-16](#)

CWM

configuring in FCIP profiles [2-29](#)

D

default gateways. See IPv4 default gateways

default networks. See IPv4 default networks

differentiated services code point. See DSCP

direct memory access devices. See DMA-bridges

DMA-bridges

displaying statistics [6-17](#)

DNS

default settings [5-8](#)

DNS hosts

displaying information [5-30](#)

DNS servers

configuring [5-7](#)

domain names

defining [5-24](#)

Domain Name System servers. See DNS servers

drivers

iSCSI [4-2](#)

DSCP

configuring [2-11](#)

dynamic initiator mode parameter

distributed with CFS [4-21](#)

dynamic iSCSI initiator

converting [4-59](#)

convert to static iSCSI

convert dynamic initiator to static [4-43](#)

dynamic mapping [4-5, 4-18](#)

dynamic mapping iSCSI

dynamic mapping iSCSI

static mapping static mapping [4-4](#)

E

ELP

verifying using Device Manager (procedure) [2-39](#)

entity status inquiry. See ESI

E ports

configuring [2-12](#)

trunking configuration [2-24](#)

ESI

non-resp threshold [4-79](#)

Ethernet MAC statistics

displaying [6-17](#)

Ethernet PortChannels

configuring [6-14](#)

description [6-9](#)

iSCSI [4-28](#)

redundancy [2-7](#)

explicit fabric logout [4-8](#)

Extended Link Protocol. See ELP

external RADIUS server

CHAP [4-73](#)

external RADIUS servers

CHAP [4-73](#)

F

fabric lock

releasing [4-69](#)

FCIP [4-1](#)

checking trunk status (procedure) [2-24](#)

compression [2-18](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- configuring [?? to 2-39](#), [?? to 2-39](#)
- configuring using FCIP Wizard [?? to 2-22](#)
- default parameters [2-19](#)
- discarding packets [2-33](#)
- enabling [2-21](#)
- Gigabit Ethernet ports [6-4](#), [7-2](#)
- high availability [2-5 to ??](#)
- IPS modules [2-2](#)
- IP storage services support [6-1](#), [6-2](#)
- link failures [2-6](#)
- MPS-14/2 module [2-2](#)
- tape acceleration [2-14 to 2-44](#)
- time stamps [2-33](#)
- VE ports [2-2](#)
- verifying ELP (procedure) [2-39](#)
- verifying interfaces (procedure) [2-39](#)
- virtual ISLs [2-2](#)
- VRRP [2-7](#)
- write acceleration [2-12](#)
- FCIP compression
 - configuring [2-37](#)
 - configuring (procedure) [2-22](#)
 - description [2-18](#)
 - displaying information [2-44](#)
- FCIP interfaces
 - configuring advanced features [?? to 2-35](#)
 - configuring peers [2-9](#)
 - configuring QoS [2-11](#)
 - creating [2-9](#)
 - displaying information [2-39](#)
 - parameters [2-5](#)
- FCIP links
 - B port interoperability mode [2-9](#)
 - configuring [2-23](#)
 - configuring peers [2-9](#)
 - configuring QoS [2-11](#)
 - description [2-3](#)
 - endpoints [2-3](#)
 - initiating IP connections [2-32](#)
 - TCP connections [2-3](#)
- FCIP listener ports
 - configuring [2-25](#)
- FCIP peers
 - configuring IP addresses [2-31](#)
- FCIP profiles
 - configuring listener ports [2-25](#)
 - configuring TCP parameters [2-26 to 2-31](#), [?? to 2-39](#)
 - creating [2-23](#)
 - description [2-4](#)
 - displaying information [2-38](#)
- FCIP tape acceleration
 - configuring [2-36](#)
 - description [2-14 to 2-18](#)
 - displaying information [2-42](#)
- FCIP TCP parameters
 - configuring buffer size [2-30](#)
 - configuring CWM [2-29](#)
 - configuring keepalive timeouts [2-26](#)
 - configuring maximum jitter [2-30](#)
 - configuring maximum retransmissions [2-27](#)
 - configuring minimum retransmit timeouts [2-26](#)
 - configuring PMTUs [2-27](#)
 - configuring SACKs [2-28](#)
 - configuring window management [2-28](#)
 - displaying [2-31](#), [2-39](#)
- FCIP write acceleration
 - configuring [2-35](#)
 - configuring (procedure) [2-22](#)
 - description [2-12](#)
 - displaying information [2-41](#)
- FCP
 - routing requests [4-4](#)
- Fibre Channel [4-1](#)
 - iSCSI targets [4-4 to 4-107](#)
- Fibre Channel interfaces
 - default settings [2-19](#), [3-4](#), [5-8](#), [6-11](#), [7-4](#), [8-13](#)
- Fibre Channel over IP. See FCIP
- Fibre Channel targets

Send documentation comments to dcnm-san-docfeedback@cisco.com

dynamic importing [4-37, 4-38](#)

dynamic mapping [4-37, 4-38](#)

Fibre Channel zoning-based access control [4-13](#)

FPSF

load balancing (example) [2-6](#)

frames

configuring MTU size [6-6, 7-6](#)

full core dumps

IPS modules [6-5](#)

G

Gigabit Ethernet

IPv4 example configuration [6-5](#)

Gigabit Ethernet interface example [4-26](#)

Gigabit Ethernet interfaces

configuring [6-4 to 6-10](#)

configuring auto-negotiation [6-6, 7-5](#)

configuring high availability [6-8 to ??](#)

configuring IPv6 addresses [8-14](#)

configuring MTU frame sizes [6-6, 7-6](#)

configuring promiscuous mode [6-6, 7-7](#)

configuring static IPv4 routing [7-8](#)

configuring VRRP [6-12](#)

default parameters [7-4](#)

displaying statistics [?? to 6-19](#)

subinterfaces [6-7, 7-3](#)

subnet requirements [6-7, 7-3](#)

verifying connectivity [7-10](#)

Gigabit Ethernet subinterfaces

configuring VLANs [7-7](#)

global authentication

parameter distributed [4-21](#)

H

HA solution example [4-25](#)

HBA port [4-7, 4-10](#)

hexadecimal fields [8-11](#)

high availability

Ethernet PortChannel [4-28](#)

Ethernet PortChannels [2-7](#)

Fibre Channel PortChannels [2-8](#)

VRRP [2-7, 4-27](#)

VRRPVRRP-based high availability [4-27](#)

ICMP

displaying statistics [6-19](#)

IPv6 [8-6](#)

ICMP packets

IPv6 header format, figure [8-7](#)

in-band management

IPFC [5-5](#)

initiators

statically mapped iSCSI [4-17](#)

interfaces

default settings [2-19, 3-4, 5-8, 6-11, 7-4, 8-13](#)

Internet Control Message Protocol. See ICMP

Internet Storage Name Service. See iSNS

IP connections

active mode [2-32](#)

initiating [2-32](#)

passive mode [2-33](#)

IPFC

configuring VSAN interfaces [5-11](#)

description [5-5](#)

enabling IPv4 routing [5-11](#)

example configuration [?? to 5-33](#)

IPS core dumps. See core dumps

IPsec

configuring with FCIP Wizard (procedure) [2-22](#)

IPS modules

configuring CDP [6-10](#)

core dumps [6-5](#)

FCIP [2-2](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- partial core dumps [6-5](#)
 - port modes [6-4, 7-2](#)
 - software upgrades [6-3](#)
 - supported features [6-1, 6-2](#)
- IPS port mode
 - description [6-4](#)
- IPS ports [4-5](#)
 - modes [7-2](#)
 - multiple connections [4-26](#)
- IP storage services
 - default parameters [6-11](#)
- IP Storage services modules. See IPS modules
- IPv4
 - configuring management interfaces [5-9](#)
 - configuring virtual routers [5-17](#)
 - default settings [7-4](#)
 - displaying statistics [7-11](#)
 - transitioning to IPv6 [8-12](#)
- IPv4 addresses
 - adding for VRRP [5-18](#)
 - configuring in VSANs [5-11](#)
 - configuring IPv4 and IPv6 protocol stacks [8-15](#)
 - IPv6 protocol stacks [8-10](#)
- IPv4 default gateways
 - configuring [5-10, 5-11](#)
 - description [5-3](#)
 - IP static routing [5-3](#)
 - static routes (tip) [5-4](#)
 - verifying configuration [5-25](#)
- IPv4 default networks
 - description [5-4](#)
- IPv4 routing
 - configuring Gigabit Ethernet interfaces [7-8](#)
 - disabling [5-11](#)
 - displaying route tables [7-10](#)
 - enabling [5-11](#)
 - verifying configuration [5-26](#)
- IPv4 static routing
 - configuring [5-12](#)
- description [5-5](#)
 - verifying configuration [5-26](#)
- IPv6
 - address types [8-3](#)
 - configuring addressing [8-11, 8-14](#)
 - configuring IPv4 and IPv6 addresses [8-15](#)
 - configuring management interfaces [5-9](#)
 - configuring neighbor discovery parameters [8-16](#)
 - configuring virtual routers [5-18](#)
 - description [?? to 8-11](#)
 - displaying information [8-20](#)
 - dual IPv4 and IPv6 protocol stack applications, figure [8-11](#)
 - dual IPv4 and IPv6 protocol stacks [8-10](#)
 - dual IPv4 and IPv6 protocol stack technique, figure [8-10](#)
 - enabling routing [8-11, 8-14](#)
 - enhancements over IPv4 [8-1](#)
 - ICMP [8-6](#)
 - neighbor discovery [8-7](#)
 - path MTU discovery [8-7](#)
 - router advertisement messages [8-9](#)
 - router discovery [8-9](#)
 - stateless autoconfiguration [8-9](#)
 - transitioning from IPv4 [8-12](#)
 - verifying basic connectivity [8-21](#)
 - verifying configuration [8-21](#)
- IPv6 addresses
 - adding for VRRP [5-19](#)
 - configuring [8-11, 8-14](#)
 - configuring IPv4 and IPv6 protocol stacks [8-15](#)
 - formats [8-2](#)
 - link-local type [8-4](#)
 - multicast type [8-5](#)
 - prefix format [8-3](#)
 - unicast type [8-3](#)
- IPv6 enhancements over IPv4 [8-1](#)
- IPv6 neighbor discovery
 - advertisement messages [8-7](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- description [8-7](#)
- neighbor solicitation message, figure [8-8](#)
- solicitation messages [8-7](#)
- IPv6 routing
 - enabling [8-11, 8-14](#)
- IPv6 static routes
 - displaying the route table [8-19](#)
- IQN
 - formats [4-5](#)
- IQNs
 - formats [4-5](#)
- ISCSI
 - enforcing access control [4-13](#)
- iSCSI
 - access control [4-11 to 4-13](#)
 - add initiator to zone database [4-47, 4-48](#)
 - advanced VSAN membershipadvanced VSAN membership [4-11](#)
 - checking for WWN conflicts [4-43](#)
 - configuring [?? to 4-29](#)
 - configuring AAA authentication [4-50, 4-51](#)
 - configuring ACLs [4-12](#)
 - configuring VRRP [4-27](#)
 - creating virtual targets [4-38](#)
 - default parameters [4-33](#)
 - discovery phase [4-13](#)
 - displaying global information [4-90](#)
 - displaying statistics [4-87](#)
 - drivers [4-2](#)
 - enabling [4-2, 4-35](#)
 - error [4-7](#)
 - Fibre Channel targets [4-4 to 4-107](#)
 - Gigabit Ethernet ports [6-4, 7-2](#)
 - GW flagiSCSI
 - gateway device [4-8](#)
 - HA with host without multi-path software [4-24](#)
 - initiator idle timeoutinitiator idle timeout
 - iSCSIinitiator idle timeout
 - configuring with Fabric Manager [4-40](#)
 - initiator name [4-52](#)
 - initiator targets [4-37](#)
 - IPS module support [6-2](#)
 - IQNs [4-6](#)
 - login redirect [4-18](#)
 - LUN mapping for targets [4-119 to 4-125](#)
 - MPS-14/2 module support [6-2](#)
 - multiple IPS ports [4-26](#)
 - PortChannel-based high availability [4-28](#)
 - PortChannel-based high availabilityEthernet
 - PortChannel-based high availability [4-28](#)
 - protocol [4-2](#)
 - requests and responses [4-4](#)
 - restrict an initiator to a specific user name for CHAP authentication [4-52](#)
 - routing [4-2](#)
 - routing modes chartrouting modes chart for iSCSI [4-16](#)
 - session creation [4-13](#)
 - session limits [4-7](#)
 - statically mapped initiators [4-17](#)
 - tables in Fabric Manager [4-44](#)
 - targets in Device Manager [4-38](#)
 - transparent initiator mode [4-7](#)
 - transparent mode initiator [4-111 to 4-116](#)
 - users with local authentication [4-52](#)
 - using iSCSI Wizard (procedure) [4-37](#)
 - VSAN membership [4-11](#)
 - VSAN membership example [4-107](#)
 - VSAN membership for iSCSI interfaces [4-11, 4-46](#)
 - zone name [4-37](#)
- iSCSI authentication
 - CHAP option [4-72](#)
 - configuring [4-13, 4-19](#)
 - configuring mechanisms [4-51](#)
 - configuring RADIUS (procedure) [4-54](#)
 - external RADIUS servers [4-73](#)
 - global override [4-50](#)
 - local authentication [4-52](#)
 - mechanisms [4-50](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- restricting on initiators [4-52](#)
- scenarios [4-71](#)
- setup guidelines [4-71](#)
- iSCSI-based access control [4-11, 4-12](#)
- iSCSI devices
 - example membership in VSANs [4-107](#)
- iscsi-gw [4-10](#)
- iSCSI high availability
 - configuring [4-23 to 4-29](#)
- iSCSI hosts
 - initiator identification [4-6](#)
- iSCSI initiators
 - assigning WWNs [4-42](#)
 - configuring dynamic IP address mapping [4-41](#)
 - configuring static IP address mapping [4-41](#)
 - displaying information [4-91 to 4-94](#)
 - displaying proxy information [4-88](#)
 - dynamic mapping [4-9](#)
 - making dynamic WWN mapping static [4-43](#)
 - proxy mode [4-9](#)
 - statically mapped (procedure) [4-41](#)
 - static mapping [4-9](#)
 - transparent mode [4-7](#)
 - WWN assignments [4-8](#)
- iSCSI interfaces
 - configuring [4-6, 4-6 to 4-55](#)
 - configuring listener ports [4-14](#)
 - configuring listener portsiSCSI
 - listener port [4-14](#)
 - configuring QoS [4-55](#)
 - configuring routing mode [4-15 to 4-55](#)
 - configuring routing modesiSCSI
 - configuring routing modesrouting modes [4-15](#)
 - configuring TCP tuning parameters [4-14](#)
 - creating [4-36](#)
 - creatingiSCSI
 - creating interfaces [4-36](#)
 - displaying information [4-86](#)
- iSCSI LUs [4-5](#)
- iSCSI protocol [4-1](#)
- iSCSI server load balancing [4-17](#)
- iSCSI Server Load Balancing. See iSLB
- iSCSI sessions
 - authentication [4-13 to 4-54](#)
 - authenticationiSCSI
 - session authenticationauthentication
 - iSCSI session [4-13](#)
 - displaying information [4-90](#)
- iSCSI targets
 - advertising [4-39](#)
 - dynamic importing [4-5](#)
 - dynamic mapping [4-5](#)
 - secondary access [4-25](#)
 - static importing [4-6](#)
 - static importingstatic mappingiSCSI targets
 - static mapping [4-6](#)
 - transparent failover [4-23, 4-23 to 4-71, ?? to 4-71](#)
- iSCSI users
 - displaying information [4-95](#)
- iSCSI virtual targets
 - displaying information [4-95](#)
- iSLB
 - activating zones [4-18, 4-61](#)
 - auto-zoning [4-22](#)
 - committing configuration changescommitting
 - configuration changes
 - iSLB [4-68](#)
 - configuration distribution [4-21 to ??, 4-67](#)
 - configuring [4-56](#)
 - configuring initiators and targets [4-18, 4-61](#)
 - configuring VRRP [4-66](#)
 - configuring with Device Manager [4-56](#)
 - configuring zones [4-18, 4-61](#)
 - default settings [4-34](#)
 - distributing configuration using CF [4-21](#)
 - dynamic initiator mapping [4-59](#)
 - enabling configuration distribution [4-67](#)
 - initiator WWN assignment [4-56](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- load balancing algorithm [4-21 to 4-110](#)
 - maximum initiators [4-32](#)
 - static initiator configuration
 - initiator configuration
 - static iSLB [4-17](#)
 - VSAN membership [4-59](#)
 - zone set activation failed [4-63](#)
 - iSib
 - default settings [4-34](#)
 - iSLB auto-zone feature [4-33](#)
 - iSLB initiators [4-18](#)
 - assigning WWNs [4-18](#)
 - configuring [?? to 4-19](#)
 - configuring IP addresses [4-57](#)
 - configuring names [4-57](#)
 - configuring static name mapping [4-58](#)
 - description [4-17](#)
 - dynamic initiator mapping [4-59](#)
 - VSAN membership [4-59](#)
 - iSLB initiator targets
 - configuring [4-61](#)
 - description [4-18, 4-61](#)
 - iSLB sessions
 - maximum per IPS port
 - maximum sessions per IPS port [4-32, 4-33](#)
 - iSLB VRRP
 - displaying information [4-95](#)
 - enabling [4-66](#)
 - iSLB with CFS distribution [4-33](#)
 - iSMS servers
 - enabling [4-78](#)
 - iSNS
 - client registration [4-31](#)
 - cloud discovery [4-105](#)
 - configuring [4-31](#)
 - configuring servers [?? to 4-31, 4-77 to ??](#)
 - description [4-29](#)
 - ESI [4-79](#)
 - iSNS client
 - description [4-30](#)
 - iSNS clients
 - creating profiles [4-75](#)
 - verifying configuration [4-97](#)
 - iSNS cloud discovery
 - CFS distribution [4-82](#)
 - description [4-31](#)
 - displaying statistics [4-105](#)
 - verifying configuration [4-104](#)
 - verifying membership [4-105](#)
 - verifying status [4-105](#)
 - iSNS profiles
 - creating [4-75](#)
 - verifying configuration [4-97](#)
 - iSNS servers
 - configuration distribution [4-78](#)
 - description [4-30](#)
 - displaying configurations [4-98 to 4-104](#)
 - enabling [4-78](#)
-
- ## J
- jitter
 - configuring estimated maximum in FCIP profiles [2-30](#)
 - jumbo frames. See MTUs
-
- ## K
- keepalive timeouts
 - configuring in FCIP profiles [2-26](#)
-
- ## L
- latency
 - forwarding [4-15](#)
 - link-local addresses
 - description [8-4](#)
 - format, figure [8-5](#)
 - link redundancy

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Ethernet PortChannel aggregation [6-9](#)
- load balancing [4-17, 4-18](#)
 - FSPF (example) [2-6](#)
 - PortChannels (example) [2-5](#)
 - weighted [4-60](#)
- load metric [4-60](#)
- lock the fabric [4-22](#)
- LUN [4-5](#)
 - trespass for storage port failover [4-71](#)
- LUN mapping [4-25](#)
 - iSCSI [4-119 to 4-125](#)
- LUNs
 - explicit access control [4-9](#)
 - mapping and assignment [4-10](#)
- LUs [4-4, 4-5](#)

M

- management interfaces
 - configuring [5-9](#)
 - configuring for IPv4 [5-9](#)
 - configuring for IPv6 [5-9](#)
 - default settings [2-19, 3-4, 5-8, 6-11, 7-4, 8-13](#)
- maximum retransmissions
 - configuring in FCIP profiles [2-27](#)
- MD5 authentication
 - VRRP [5-22](#)
- merge status conflictsiSLB
 - merge status conflictsCFS
 - merge status conflicts [4-23](#)
- mgmt0 interfaces
 - configuring IPv4 addresses [5-9](#)
 - configuring IPv6 addresses [5-9](#)
 - default settings [2-19, 3-4, 5-8, 6-11, 7-4, 8-13](#)
 - local IPv4 routing [5-4](#)
- minimum retransmit timeouts
 - configuring in FCIP profiles [2-26](#)
- MPS-14/2 modules [4-1, 4-2, 4-3, 4-10, 4-13, 4-36](#)
 - FCIP [2-2](#)

- port modes [6-4, 7-2](#)
- software upgrades [6-4](#)
- supported features [6-1, 6-2](#)

- MTU frame sizes
 - configuring Gigabit Ethernet interfaces [6-6](#)

- MTUs
 - configuring frame sizes [7-6](#)
 - configuring size
 - path discovery for IPv6 [8-7](#)

- multicast addresses
 - IPv6 alternative to broadcast addresses [8-6](#)
 - IPv6 format, figure [8-5](#)
 - IPv6 solicited-node format, figure [8-6](#)

- multi-path software example [4-24](#)

- multiple VSANs
 - configuring [5-14](#)

- Multiprotocol Services modules. See MPS-14/2 modules

- mutual CHAP authentication
 - configuring for iSCSI [4-53](#)
 - configuring for iSLB [4-65](#)
 - configuring for iSLBI [4-65](#)

N

- neighbor discovery
 - configuring parameters [8-16](#)
 - verifying configuration [8-19](#)

- None authentication [4-13](#)

- NTP
 - time-stamp option [2-33](#)

O

- overlay VSANs
 - configuring [5-12](#)
 - description [5-5](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

P

packets

discarding in FCIP [2-33](#)

pass-thru routing mode [4-15, 4-16](#)

path MTUs. See PMTUs

PDU [4-15](#)

PMTUs

configuring in FCIP profiles [2-27](#)

PortChannel

interfaces [4-39](#)

subinterfaces [4-39](#)

PortChannels

configuring for FCIP high availability [2-5](#)

IQN formats [4-5](#)

load balancing (example) [2-5](#)

member combinations [6-9](#)

redundancy [2-8](#)

port modes

IPS [6-4, 7-2](#)

ports

virtual E [2-2](#)

promiscuous mode

configuring Gigabit Ethernet interfaces [6-6, 7-7](#)

protocol [4-1](#)

protocols

VRRP [4-5](#)

proxy initiator

configuring iSCSI

configuring proxy initiator [4-44, 4-45](#)

proxy initiator mode [4-7, 4-12](#)

configuring [4-10](#)

zoning [4-45](#)

proxy initiator mode iSCSI

proxy initiator mode [4-9](#)

Pv6 address formats [8-11](#)

pWWNs

converting dynamic to static [4-43](#)

Q

QoS

DSCP value [2-11](#)

QoS values

configuring [4-55](#)

R

RADIUS [4-74](#)

AAA authentication [4-13, 4-19](#)

configuring an iSCSI RADIUS server iSCSI

configuring a RADIUS server [4-54](#)

redundancy

Ethernet PortChannels [2-7, 2-8](#)

Fibre Channel PortChannels [2-8](#)

VRRP [2-7](#)

router discovery

IPv6 [8-9](#)

RSCNs [4-40](#)

S

SACKs

configuring in FCIP profiles [2-28](#)

SAN extension tuner

assigning SCSI read/write commands [3-8, 3-10](#)

configuring [3-3](#)

configuring data patterns [3-11](#)

configuring nWWNs [3-7](#)

configuring virtual N ports [3-7](#)

data patterns [3-4](#)

initialization [3-7](#)

tuning guidelines [3-2](#)

verifying configuration [3-12](#)

SCSI

routing requests [4-2](#)

security parameter index. See SPI

selective acknowledgments. See SACKs

Send documentation comments to dcnm-san-docfeedback@cisco.com

SPI

configuring virtual routers [5-22](#)

statically imported iSCSI targets [4-25](#)

static iSLB initiator

converting [4-59](#)

static mapped iSCSI target

static mapped target [4-13](#)

static mapping [4-18](#)

static WWN mapping [4-12](#)

store-and-forward routing mode [4-15, 4-16](#)

subnet masks

configuring IPv4 routes [5-12](#)

subnets

requirements [6-7, 7-3](#)

switch management

in-band [5-5](#)

switchovers

VRRP [2-7](#)

T

TACACS+

AAA authentication [4-19](#)

target discovery [4-31](#)

TCP connections

FCIP profiles [2-4](#)

TCP parameters

configuring in FCIP profiles [2-26 to 2-31, ?? to 2-39](#)

TCP statistics

displaying [6-18](#)

TCP tuning parameters [4-14](#)

transient failure [4-40](#)

transparent initiator mode [4-7](#)

transparent initiator mode

transparent initiator mode [4-9](#)

troubleshooting

CTC [2-24](#)

trunking

link state [2-24](#)

trunking mode

FCIP interface [2-5](#)

V

VE ports

description [2-2](#)

FCIP [2-2](#)

virtual E ports. See VE ports

virtual Fibre Channel host [4-3](#)

virtual ISLs

description [2-2](#)

Virtual LANs. See VLANs

virtual LANs. See VLANs

virtual router IDs. See VR IDs

Virtual Router Redundancy Protocol. See VRRP

Virtual Router Redundancy Protocol

Virtual Router Redundancy [4-17](#)

virtual routers

adding [5-17](#)

adding primary IP addresses [5-18](#)

authentication [5-22](#)

configuring for IPv4 [5-17](#)

configuring for IPv6 [5-18](#)

default settings [5-8](#)

deleting [5-17](#)

initiating [5-17](#)

setting priorities [5-19](#)

VLANs

configuring on Gigabit Ethernet subinterfaces [7-7](#)

description [6-6, 7-2](#)

VR IDs

configuring for IPv4 [5-17](#)

configuring for IPv6 [5-17](#)

description [5-6](#)

mapping [5-6](#)

VRRP [4-17](#)

algorithm for selecting Gigabit Ethernet interfaces [4-21 to 4-110](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- backup switches [5-6](#)
- clearing statistics [5-29](#)
- configuring advertisement time intervals [5-20](#)
- configuring for Gigabit Ethernet interfaces [6-12](#)
- configuring for iSLB [4-66](#)
- configuring virtual routers [5-17](#)
- configuring VR IDs for IPv4 [5-17](#)
- configuring VR IDs for IPv6 [5-17](#)
- default settings [5-8](#)
- description [5-5, 6-8](#)
- displaying information [5-27 to 5-29](#)
- displaying statistics [5-29](#)
- group members [6-8](#)
- initiating virtual routers [5-17](#)
- IQN formats [4-5](#)
- iSCSI parameter change impact [4-21](#)
- iSLB [4-19 to 4-96](#)
- master switches [5-6](#)
- MD5 authentication [5-22](#)
- primary IP address [5-18](#)
- priority preemption [5-21](#)
- security authentication [5-22](#)
- setting priorities [5-19](#)
- setting priority [5-19](#)
- simple text authentication [5-22](#)

VRRP group [4-46](#)

VRRP–I f iSCSI login redirect [4-18](#)

VSAN interfaces

- configuring IPv4 addresses [5-11](#)

- verifying configuration [5-26](#)

VSANs

- configuring multiple IPv4 subnets [5-14](#)

- example membership for iSCSI devices [4-107](#)

- gateway switches [5-4](#)

- IPv4 static routing [5-5](#)

- iSLB [4-59](#)

- iSLB initiators [4-59](#)

- traffic routing between [5-1](#)

- VRRP [5-6](#)

W

window management

- configuring in FCIP profiles [2-28](#)

WWNs

- static binding [4-10](#)

Z

zones

- configuring and activating for iSLB [4-18, 4-61](#)

- iSLB [4-18, 4-61](#)

zoning based access control

- configuring for iSCSI [4-11](#)

- configuring for iSCSIiSCSI

- configuring zoning based access control [4-11](#)