



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



## **Cisco Fabric Manager IP Services Configuration Guide Release 5.0(1a)**

Cisco Fabric Manager Release 5.0(1a)  
February 2010

### **Americas Headquarters**

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

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.

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (1002R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.



## CONTENTS

### New and Changed Information xi

### Preface xiii

Audience xiii

Organization xiii

Document Conventions xiv

Related Documentation xiv

Release Notes xv

Regulatory Compliance and Safety Information xv

Compatibility Information xv

Hardware Installation xv

Software Installation and Upgrade xv

Cisco NX-OS xv

Cisco Fabric Manager xvi

Command-Line Interface xvi

Intelligent Storage Networking Services Configuration Guides xvi

Troubleshooting and Reference xvii

Obtaining Documentation and Submitting a Service Request xvii

---

## CHAPTER 1

### IP Services Overview 1-1

FCIP 1-1

SAN Extension Tuner 1-2

iSCSI 1-2

IP Services 1-2

IP Storage 1-2

IPv4 and IPv6 1-2

---

## CHAPTER 2

### Configuring FCIP 2-1

About FCIP 2-1

FCIP Concepts 2-2

FCIP and VE Ports 2-2

FCIP Links 2-3

FCIP Profiles 2-4

FCIP Interfaces 2-4

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

FCIP High-Availability Solutions	2-4
Fibre Channel PortChannels	2-5
FSPF	2-5
VRRP	2-6
Ethernet PortChannels	2-6
Ethernet PortChannels and Fibre Channel PortChannels	2-7
Configuring FCIP	2-7
Enabling FCIP	2-8
Using the FCIP Wizard	2-8
Basic FCIP Configuration	2-15
Creating FCIP Profiles	2-15
Creating FCIP Links	2-16
Verifying Interfaces and Extended Link Protocol	2-16
Checking Trunk Status	2-17
Launching Cisco Transport Controller	2-17
Advanced FCIP Profile Configuration	2-17
Configuring TCP Parameters	2-17
Advanced FCIP Interface Configuration	2-20
Configuring Peers	2-20
Assigning a Peer IP Address	2-21
Configuring Active Connections	2-22
Enabling Time Stamp Control	2-22
FCIP B Port Interoperability Mode	2-22
Quality of Service	2-25
Configuring E Ports	2-25
Advanced FCIP Features	2-26
FCIP Write Acceleration	2-26
Configuring FCIP Write Acceleration	2-28
FCIP Tape Acceleration	2-28
Configuring FCIP Tape Acceleration	2-33
FCIP Compression	2-33
Default Settings	2-34

## CHAPTER 3

### Configuring the SAN Extension Tuner 3-1

About the SAN Extension Tuner	3-1
SAN Extension Tuner Setup	3-2
Data Pattern	3-3
License Prerequisites	3-3
Configuring the SAN Extension Tuner	3-3

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

Tuning the FCIP Link	3-4
Using the SAN Extension Tuner Wizard	3-4
Default Settings	3-6

## CHAPTER 4

<b>Configuring iSCSI</b>	<b>4-1</b>
About iSCSI	4-1
About iSCSI Configuration Limits	4-4
Configuring iSCSI	4-4
Enabling iSCSI	4-4
Creating iSCSI Interfaces	4-6
Using the iSCSI Wizard	4-6
Presenting Fibre Channel Targets as iSCSI Targets	4-8
Dynamic Mapping	4-9
Static Mapping	4-11
iSCSI Virtual Target Configuration Examples	4-13
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts	4-15
Initiator Identification	4-15
Initiator Presentation Modes	4-15
VSAN Membership for iSCSI	4-23
Example of VSAN Membership for iSCSI Devices	4-24
Advanced VSAN Membership for iSCSI Hosts	4-25
iSCSI Access Control	4-25
Fibre Channel Zoning-Based Access Control	4-26
iSCSI-Based Access Control	4-27
Enforcing Access Control	4-28
iSCSI Session Authentication	4-29
Configuring Authentication Mechanism	4-30
Configuring Local Authentication	4-31
Restricting iSCSI Initiator Authentication	4-31
Configuring Mutual CHAP Authentication	4-31
Configuring an iSCSI RADIUS Server	4-32
iSCSI Immediate Data and Unsolicited Data Features	4-32
iSCSI Interface Advanced Features	4-33
iSCSI Listener Port	4-33
TCP Tuning Parameters	4-33
Setting QoS Values	4-33
iSCSI Routing Modes	4-34
Configuring iSLB	4-36
About iSLB Configuration Limits	4-37

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

iSLB Configuration Prerequisites	4-38
About iSLB Initiators	4-38
Configuring iSLB Using Device Manager	4-38
Configuring iSLB Initiators	4-40
Assigning WWNs to iSLB Initiators	4-41
Making the Dynamic iSLB Initiator WWN Mapping Static	4-41
Assigning VSAN Membership for iSLB Initiators	4-41
Configuring Metric for Load Balancing	4-42
Configuring iSLB Initiator Targets	4-42
Configuring and Activating Zones for iSLB Initiators and Initiator Targets	4-43
Configuring iSLB Session Authentication	4-44
About Load Balancing Using VRRP	4-44
Changing iSCSI Interface Parameters and the Impact on Load Balancing	4-46
VRRP Load Balancing Algorithm For Selecting Gigabit Ethernet Interfaces	4-46
Configuring Load Balancing Using VRRP	4-46
About iSLB Configuration Distribution Using CFS	4-47
Distributing the iSLB Configuration Using CFS	4-48
Enabling iSLB Configuration Distribution	4-48
Locking the Fabric	4-49
Committing Changes to the Fabric	4-49
Discarding Pending Changes	4-49
Clearing a Fabric Lock	4-50
CFS Merge Process	4-50
iSLB CFS Merge Status Conflicts	4-50
iSCSI High Availability	4-51
Transparent Target Failover	4-51
iSCSI High Availability with Host Running Multi-Path Software	4-51
iSCSI HA with Host Not Having Any Multi-Path Software	4-52
LUN Trespass for Storage Port Failover	4-54
Multiple IPS Ports Connected to the Same IP Network	4-54
VRRP-Based High Availability	4-55
Ethernet PortChannel-Based High Availability	4-56
iSCSI Authentication Setup Guidelines and Scenarios	4-57
Configuring No Authentication	4-57
Configuring CHAP with Local Password Database	4-58
Configuring CHAP with External RADIUS Server	4-58
iSCSI Transparent Mode Initiator	4-59
Target Storage Device Requiring LUN Mapping	4-63
iSNS	4-68

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

About iSNS Client Functionality	4-68
Creating an iSNS Client Profile	4-69
About iSNS Server Functionality	4-70
Example Scenario	4-71
Configuring iSNS Servers	4-72
Enabling the iSNS Server	4-72
iSNS Configuration Distribution	4-72
Configuring the ESI Retry Count	4-73
Configuring the Registration Period	4-73
iSNS Client Registration and Deregistration	4-73
Target Discovery	4-73
iSNS Cloud Discovery	4-74
About Cloud Discovery	4-74
Configuring iSNS Cloud Discovery	4-75
Enabling iSNS Cloud Discovery	4-75
Initiating On-Demand iSNS Cloud Discovery	4-75
Configuring Automatic iSNS Cloud Discovery	4-76
Configuring iSNS Cloud Discovery Distribution	4-76
Default Settings	4-76

## CHAPTER 5

### Configuring IP Services 5-1

Traffic Management Services	5-2
Management Interface Configuration	5-2
Default Gateway	5-3
About the Default Gateway	5-3
Configuring the Default Gateway	5-3
IPv4 Default Network Configuration	5-6
IPFC	5-7
IPFC Configuration Guidelines	5-7
IPv4 Static Routes	5-7
Overlay VSANs	5-7
About Overlay VSANs	5-8
Configuring Overlay VSANs	5-8
Configuring Multiple VSANs	5-9
Virtual Router Redundancy Protocol	5-10
About VRRP	5-11
Configuring VRRP	5-12
Adding and Deleting a Virtual Router	5-12

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

Virtual Router Initiation	5-13
Adding Virtual Router IP Addresses	5-13
Setting the Priority for the Virtual Router	5-13
Setting the Time Interval for Advertisement Packets	5-13
Configuring or Enabling Priority Preemption	5-13
Setting Virtual Router Authentication	5-14
Tracking the Interface Priority	5-14
DNS Server Configuration	5-14
Default Settings	5-15

## CHAPTER 6

### Configuring IP Storage 6-1

Services Modules	6-1
Module Status Verification	6-2
IPS Module Upgrade	6-3
MPS-14/2 Module Upgrade	6-3
Supported Hardware	6-3
Configuring Gigabit Ethernet Interfaces for IPv4	6-4
Basic Gigabit Ethernet Configuration	6-4
Configuring Interface Descriptions	6-5
Configuring Beacon Mode	6-5
Configuring Autonegotiation	6-5
Configuring the MTU Frame Size	6-5
Configuring Promiscuous Mode	6-6
About VLANs for Gigabit Ethernet	6-6
Interface Subnet Requirements	6-6
Verifying Gigabit Ethernet Connectivity	6-7
Gigabit Ethernet IPv4-ACL Guidelines	6-7
Configuring Gigabit Ethernet High Availability	6-8
VRRP for iSCSI and FCIP Services	6-8
Configuring VRRP for Gigabit Ethernet Interfaces	6-9
About Ethernet PortChannel Aggregation	6-9
Configuring Ethernet PortChannels	6-10
Configuring CDP	6-10
Default Settings	6-10

## CHAPTER 7

### Configuring IPv4 for Gigabit Ethernet Interfaces 7-1

About IPv4	7-1
Basic Gigabit Ethernet Configuration for IPv4	7-2



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

Configuring Interface Descriptions	7-3
Configuring Beacon Mode	7-3
Configuring Autonegotiation	7-3
Configuring the MTU Frame Size	7-3
Configuring Promiscuous Mode	7-4
VLANs	7-4
About VLANs for Gigabit Ethernet	7-4
Configuring the VLAN Subinterface	7-5
Interface Subnet Requirements	7-5
IPv4-ACLs	7-5
Gigabit Ethernet IPv4-ACL Guidelines	7-6
Default Settings	7-6

## CHAPTER 8

### **Configuring IPv6 for Gigabit Ethernet Interfaces 8-1**

About IPv6	8-1
Extended IPv6 Address Space for Unique Addresses	8-2
IPv6 Address Formats	8-2
IPv6 Address Prefix Format	8-3
IPv6 Address Type: Unicast	8-3
Global Addresses	8-3
Link-Local Address	8-4
IPv6 Address Type: Multicast	8-5
ICMP for IPv6	8-6
Path MTU Discovery for IPv6	8-7
IPv6 Neighbor Discovery	8-7
IPv6 Neighbor Solicitation and Advertisement Messages	8-7
Router Discovery	8-9
IPv6 Stateless Autoconfiguration	8-9
Dual IPv4 and IPv6 Protocol Stacks	8-10
Configuring Basic Connectivity for IPv6	8-11
Configuring IPv6 Addressing and Enabling IPv6 Routing	8-11
Configuring IPv4 and IPv6 Protocol Addresses	8-13
Configuring IPv6 Static Routes	8-13
Configuring a IPv6 Static Route	8-13
Gigabit Ethernet IPv6-ACL Guidelines	8-14
Transitioning from IPv4 to IPv6	8-15
Default Settings	8-15

## INDEX

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



## New and Changed Information

As of Cisco MDS NX-OS Release 5.0(1a), software configuration information is available in new feature-specific configuration guides for the following information:

- System management
- Interfaces
- Fabric
- Quality of service
- Security
- IP services
- High availability and redundancy

The information in these new guides previously existed in the *Cisco MDS 9000 Family CLI Configuration Guide* and in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. Those configuration guides remain available on Cisco.com and should be used for all software releases prior to MDS FM Release 5.0(1a). Each guide addresses the features introduced in or available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

For a complete list of document titles, see the list of Related Documentation in the “Preface.”

To find additional information about Cisco MDS FM Release 5.0(1a), see the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:

[http://www.cisco.com/en/US/products/ps5989/prod\\_release\\_notes\\_list.htm](http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.htm)

### 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*.

[Table 1](#) lists the New and Changed features for this guide, starting with Cisco Fabric Manager IP Services Release 5.0(1a)

**Table 1**      **New and Changed Features for Cisco Fabric Manager IP Services Release 5.0(1a)**

Feature	New or Changed Topics	Changed in Release	Where Documented
Configuring FCIP	FCIP Compression	5.0(1a)	<a href="#">“FCIP Compression” section on page 33</a>

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



## Preface

---

This preface describes the audience, organization, and conventions of the *Cisco Fabric Manager IP Services Configuration Guide*. It also provides information on how to obtain related documentation.

## Audience

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

## Organization

This guide is organized as follows:

Chapter	Title	Description
<a href="#">Chapter 1</a>	<a href="#">IP Services Overview</a>	Provides an overview of the Intelligent Storage Services supported by the Cisco MDS 9000 NX-OS software.
<a href="#">Chapter 2</a>	<a href="#">Configuring FCIP</a>	Describes how the switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.
<a href="#">Chapter 3</a>	<a href="#">Configuring the SAN Extension Tuner</a>	Explains the SAN extension tuner (SET) feature that optimizes FCIP performance.
<a href="#">Chapter 4</a>	<a href="#">Configuring iSCSI</a>	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.
<a href="#">Chapter 5</a>	<a href="#">Configuring IP Services</a>	Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.
<a href="#">Chapter 6</a>	<a href="#">Configuring IP Storage</a>	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.

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

Chapter	Title	Description
<a href="#">Chapter 7</a>	<a href="#">Configuring IPv4 for Gigabit Ethernet Interfaces</a>	Describes the IPv4 protocol support provided by Cisco MDS 9000 Family switches.
<a href="#">Chapter 8</a>	<a href="#">Configuring IPv6 for Gigabit Ethernet Interfaces</a>	Describes the IPv6 protocol support provided by Cisco MDS 9000 Family switches.

## Document Conventions

Command descriptions use these conventions:

<b>boldface font</b>	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:

screen font	Terminal sessions and information the switch displays are in screen font.
<b>boldface screen font</b>	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](http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocator.htm)

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

## 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 Storage Services Interface Images*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Release Notes for Cisco MDS 9000 Family Fabric Manager*

## 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 NX-OS Release Compatibility Matrix for Storage Service Interface Images*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*
- *Cisco MDS NX-OS Release Compatibility Matrix for IBM SAN Volume Controller Software for Cisco MDS 9000*
- *Cisco MDS SAN-OS Release Compatibility Matrix for VERITAS Storage Foundation for Networks Software*

## 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 Release 4.1(x) and SAN-OS 3(x) Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Storage Services Interface Image Install and Upgrade Guide*
- *Cisco MDS 9000 Family Storage Services Module Software Installation and Upgrade Guide*

## Cisco NX-OS

- *Cisco MDS 9000 Family NX-OS Licensing Guide*
- *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*

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

- *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*
- *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*

## Cisco Fabric Manager

- *Cisco Fabric Manager Fundamentals Configuration Guide*
- *Cisco Fabric Manager System Management Configuration Guide*
- *Cisco Fabric Manager Interfaces Configuration Guide*
- *Cisco Fabric Manager Fabric Configuration Guide*
- *Cisco Fabric Manager Quality of Service Configuration Guide*
- *Cisco Fabric Manager Security Configuration Guide*
- *Cisco Fabric Manager IP Services Configuration Guide*
- *Cisco Fabric Manager Intelligent Storage Services Configuration Guide*
- *Cisco Fabric Manager High Availability and Redundancy Configuration Guide*
- *Cisco Fabric Manager Inter-VSAN Routing Configuration Guide*
- *Cisco Fabric Manager Online Help*
- *Cisco Fabric Manager Web Services Online Help*

## Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

## Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 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*
- *Cisco MDS 9000 Family Secure Erase Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*



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

## Troubleshooting and Reference

- *Cisco NX-OS System Messages Reference*
- *Cisco MDS 9000 Family NX-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco MDS 9000 Family NX-OS SMI-S Programming Reference*
- *Cisco MDS 9000 Family Fabric Manager Server Database Schema*

## 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***



# CHAPTER 1

## 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 [fm-docfeedback@cisco.com](mailto:fm-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 Fabric Manager 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 [fm-docfeedback@cisco.com](mailto:fm-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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***



## CHAPTER 2

# 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 sections:

- [About FCIP, page 2-1](#)
- [Configuring FCIP, page 2-7](#)
- [Using the FCIP Wizard, page 2-8](#)
- [Default Settings, page 2-34](#)

## 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.

***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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 2-4](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 2-7](#)

## FCIP Concepts

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

- [FCIP and VE Ports, page 2-2](#)
- [FCIP Links, page 2-3](#)
- [FCIP Profiles, page 2-4](#)
- [FCIP Interfaces, page 2-4](#)

## 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).

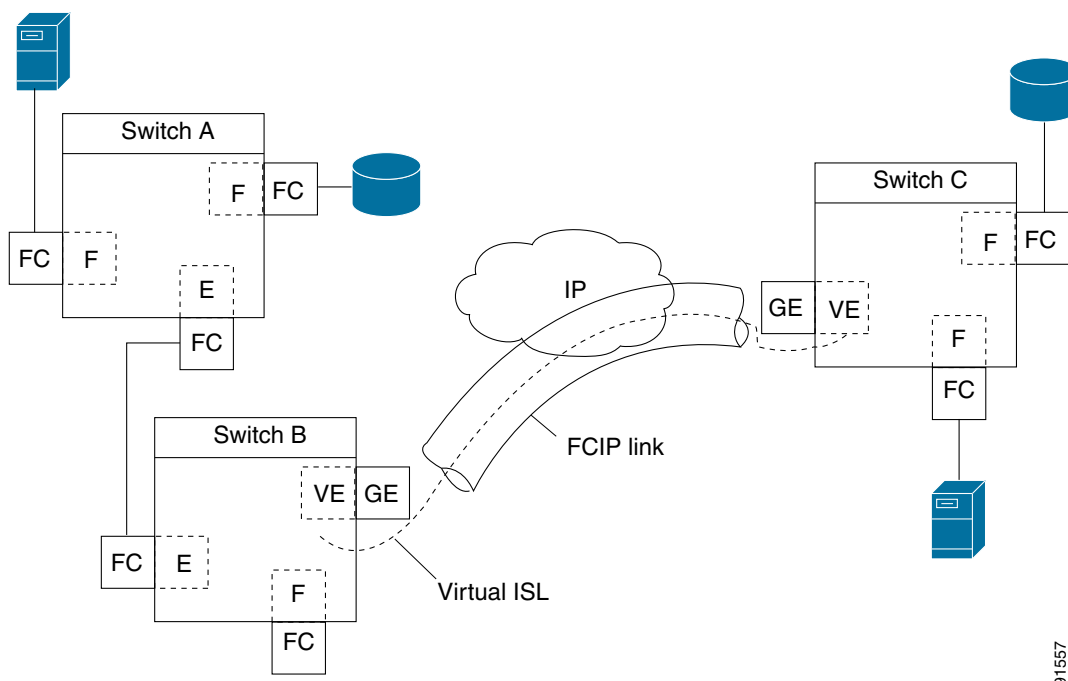
FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 2-2 FCIP Links and Virtual ISLs**



See the “Configuring E Ports” section on page 2-25 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.

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

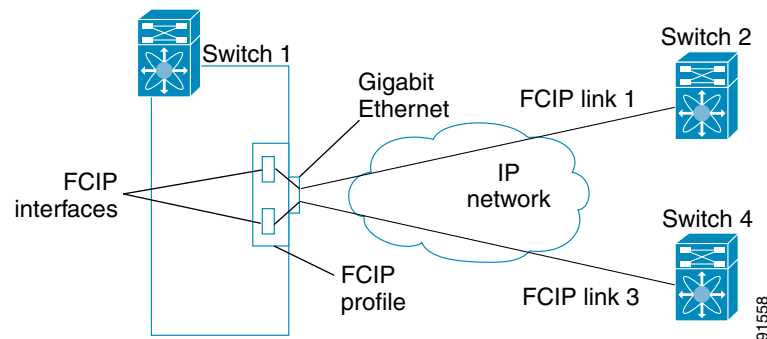
## FCIP Profiles

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 2-3](#)).

**Figure 2-3** *FCIP Profile and FCIP Links*



## FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E 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 2-5
- [VRRP](#), page 2-6
- [Ethernet PortChannels](#), page 2-6

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

## Fibre Channel PortChannels

Figure 2-4 provides an example of a PortChannel-based load-balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

**Figure 2-4 PortChannel-Based Load Balancing**




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

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

## FSPF

Figure 2-5 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 2-5 FSPF-Based Load Balancing**



The following characteristics set FSPF solutions apart from other solutions:


- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP link.

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

## VRRP

Figure 2-6 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 2-6 VRRP-Based High Availability**




The following characteristics set VRRP solutions apart from other solutions:

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

## Ethernet PortChannels

Figure 2-7 displays an Ethernet PortChannel-based high-availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

**Figure 2-7 Ethernet PortChannel-Based High Availability**



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

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

**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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 Gigabit Ethernet High Availability”](#) 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 2-8](#)).

**Figure 2-8 PortChannels at the Fibre Channel and Ethernet Levels**



To configure Fibre Channel PortChannels, see the *Cisco Fabric Manager Interfaces Configuration Guide*.

To configure Ethernet PortChannels, see the *Cisco Fabric Manager High Availability and Redundancy Configuration Guide*.

## Configuring FCIP

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

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

- [Enabling FCIP, page 2-8](#)
- [Basic FCIP Configuration, page 2-15](#)
- [Verifying Interfaces and Extended Link Protocol, page 2-16](#)
- [Checking Trunk Status, page 2-17](#)
- [Advanced FCIP Profile Configuration, page 2-17](#)
- [Advanced FCIP Interface Configuration, page 2-20](#)
- [Configuring E Ports, page 2-25](#)
- [Configuring E Ports, page 2-25](#)
- [Advanced FCIP Features, page 2-26](#)

## Enabling 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). 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 for the FCIP feature are available only when FCIP is enabled on a switch.

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.

## Using the FCIP Wizard



### 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 Fabric Manager, 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 the 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:

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

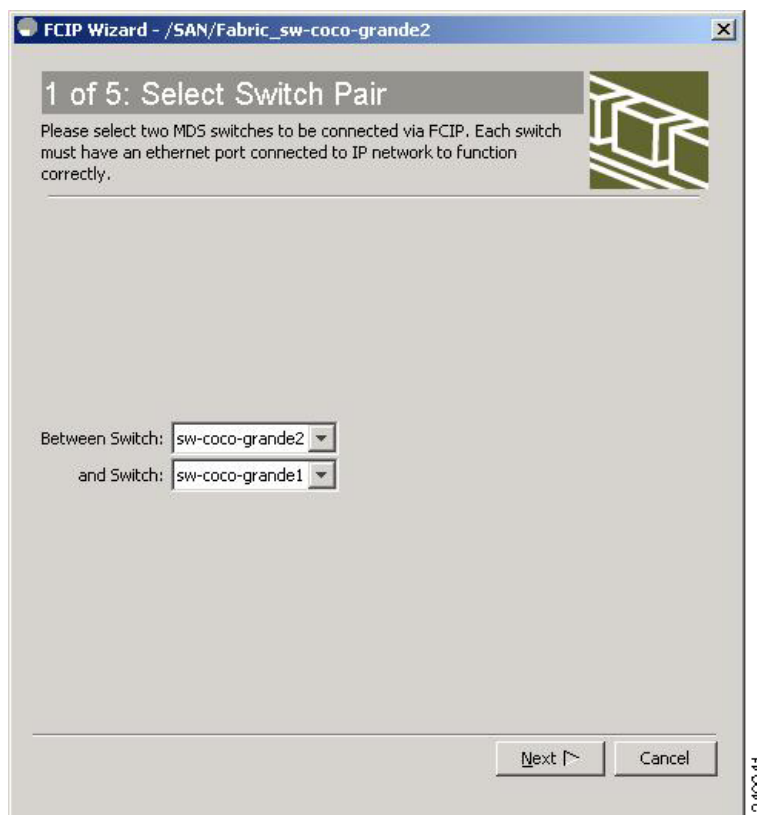
- Step 1** Click the **FCIP Wizard** icon in the Fabric Manager toolbar (See [Figure 2-9](#)).

**Figure 2-9 FCIP Wizard**



You see the switch selections as shown in [Figure 2-10](#).

**Figure 2-10 Switch Selections**



- 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 [Figure 2-11](#)). See the *Cisco Fabric Manager Security Configuration Guide* for information on IPsec and IKE.

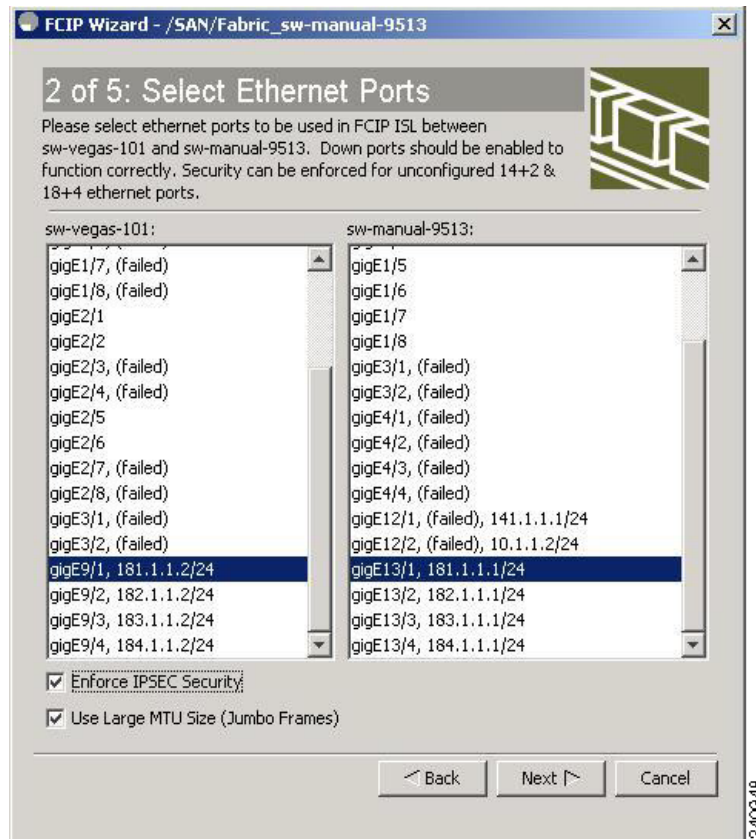
Check the **Use Large MTU Size (Jumbo Frames)** option to use jumbo size frames of 2300. Since Fibre Channel frames are 2112, we recommended that you use this option. If you uncheck the box, the FCIP Wizard does not set the MTU size, and the default value of 1500 is set.

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

**Note**

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

**Figure 2-11** Enabling IPsec on an FCIP link



**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 (see [Figure 2-12](#)).



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

**Figure 2-12 Specify IP Address/Route**

**FCIP Wizard - /SAN/Fabric\_sw-manual-9513**

**3 of 5: Specify IP Address/Route**

Please supply Ethernet Port IP Address. Specify Route if the Port addresses are in different subnet.  
Note: The changes to IP Address and IP Route Addition will be applied on pressing the Next button.

**Switch sw-vegas-101 (gigE9/1)**

IP Address/Mask:  e.g. 10.1.1.1/24  
Dest/Mask:  e.g. 10.1.0.0/16  
☐ Add IP Route: Gateway:  e.g. 11.1.1.1  
Metric:  0,,32766

**Switch sw-manual-9513 (gigE13/1)**

IP Address/Mask:  e.g. 10.1.1.1/24  
Dest/Mask:  e.g. 10.1.0.0/16  
☐ Add IP Route: Gateway:  e.g. 11.1.1.1  
Metric:  0,,32766

240976

**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 (see [Figure 2-13](#)).

You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.

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

**Figure 2-13 Specifying Tunnel Properties**

FCIP Wizard - /SAN/Fabric

**4 of 5: Specify Tunnel Properties**

Please supply the following parameters to tune the TCP connections. If Write Acceleration is enabled, ensure that flows will not load balanced across multiple ISLs.

Max Bandwidth:   1..1000 Mb

Min Bandwidth:  Mb

Estimated RTT (RoundTrip Time):   0..300000 us

☐ Write Acceleration

☐ Enable Optimum Compression

☐ Enable XRC Emulator

274886

- 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 2-26.
- Step 10** Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the “[FCIP Compression](#)” section on page 2-33.
- 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 *Cisco Fabric Manager Fabric Configuration Guide*.
- Step 12** Click **Next**.
- Step 13** Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link (see [Figure 2-14](#)).



**Note**

If FICON is enabled/FICON VSAN is present on both the switches, the [Figure 2-16](#) is displayed, otherwise [Figure 2-17](#) is displayed.

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

**Figure 2-14 Create FCIP ISL**

**5 of 5: Create FCIP ISL**

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.  
NOTE: the FCIP link may take time to appear in map.

Between Switch sw-vegas-101 (fcip13 over gigE9/1[Secure])  
And Switch sw-manual-9513 (fcip13 over gigE13/1[Secure])

**Attributes**

Port VSAN: 1 1..4093

Trunk Mode: ☒ nonTrunk ☐ trunk ☐ auto

Back Finish Cancel

240945

**Figure 2-15 Enter FICON Port Address**

**5 of 5: Create FCIP ISL**

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.  
NOTE: the FCIP link may take time to appear in map.

Between Switch v-190 (fcip6 over gigE2/1)  
FICON Port Address 0xe0..0xf9

And Switch v-172.22.31.184 (fcip6 over gigE3/1)  
FICON Port Address 0xe0..0xf9

**Attributes**

Port VSAN: 1 1..4093

Trunk Mode: ☒ nonTrunk ☐ trunk ☐ auto

FICON Tape Acceleration: ☐ VSAN List:

Back Finish Cancel

240947

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

**Figure 2-16 Create FCIP ISL**

**FCIP Wizard - /SAN/Fabric\_sw-manual-9513**

**5 of 5: Create FCIP ISL**

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.  
NOTE: the FCIP link may take time to appear in map.

Between Switch **sw-vegas-101 (fcip13 over gigE9/1[Secure])**

And Switch **sw-manual-9513 (fcip13 over gigE13/1[Secure])**

**Attributes**

Port VSAN:  1..4093

Trunk Mode: ☒ nonTrunk ☐ trunk ☐ auto

240946

**Figure 2-17 Enter FICON Port Address**

**FCIP Wizard - /SAN/Fabric\_c-186**

**5 of 5: Create FCIP ISL**

Please supply following parameters to create a FCIP tunnel. Specify Port VSAN for nontrunk/auto and allowed VSAN list for Trunk tunnel.  
NOTE: the FCIP link may take time to appear in map.

Between Switch **v-190 (fcip6 over gigE2/1)**

FICON Port Address  0xe0..0xf9

And Switch **v-172.22.31.184 (fcip6 over gigE3/1)**

FICON Port Address  0xe0..0xf9

**Attributes**

Port VSAN:  1..4093

Trunk Mode: ☒ nonTrunk ☐ trunk ☐ auto

FICON Tape Acceleration: ☐ VSAN List:

240947

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

**Step 14** Click **Finish** to create this FCIP link.

## Basic FCIP Configuration

Once you have created FCIP links using the FCIP wizard, you may need to modify parameters for these links. This 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 configure 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

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 2-18](#) shows an example configuration.

**Figure 2-18** Assigning Profiles to Each 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 Fabric Manager, 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 Fabric Manager or the **Create** button on Device Manager to add a new profile.
  - 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 Fabric Manager Online Help for explanations of these fields.

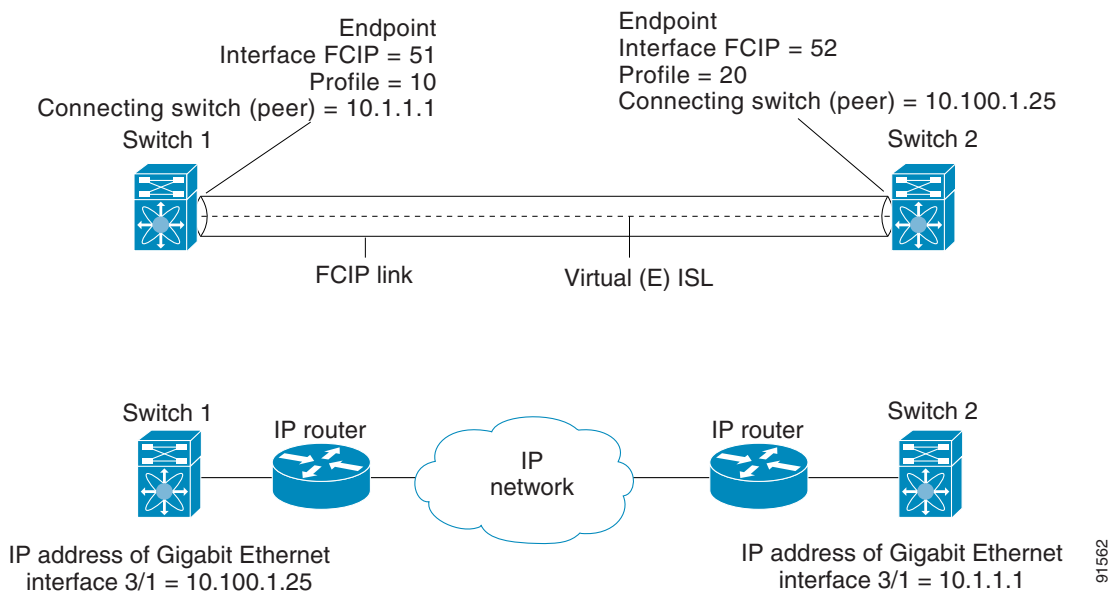
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

- 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.

## Creating FCIP Links

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 2-19](#)).

**Figure 2-19** Assigning Profiles to Each Gigabit Ethernet Interface



## Verifying Interfaces and Extended Link Protocol

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All 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.

***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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.
- 

## 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.

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.

To launch CTC using Fabric Manager, 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.
  - Step 4** Click **OK**.
- 

## Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

This section includes the following topics:

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

- [Configuring TCP Parameters, page 2-18](#)

## 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:

- [Minimum Retransmit Timeout, page 2-18](#)
- [Keepalive Timeout, page 2-18](#)
- [Maximum Retransmissions, page 2-18](#)
- [Path MTUs, page 2-19](#)
- [Selective Acknowledgments, page 2-19](#)
- [Window Management, page 2-19](#)
- [Monitoring Congestion, page 2-19](#)
- [Estimating Maximum Jitter, page 2-20](#)
- [Buffer Size, page 2-20](#)

### Minimum Retransmit Timeout

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

### Keepalive Timeout

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

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. The keepalive timeout feature 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.

### Maximum Retransmissions

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



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

## 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.

## 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.

## 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.

## 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.

## ***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***

- 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.

## **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.

## **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.

## **Advanced FCIP Interface Configuration**

This section describes the options you can configure on an FCIP interface to establish connection to a peer and includes the following topics:

- [Configuring Peers, page 2-21](#)
- [Assigning a Peer IP Address, page 2-21](#)
- [Configuring Active Connections, page 2-22](#)

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

- [Enabling Time Stamp Control, page 2-22](#)
- [FCIP B Port Interoperability Mode, page 2-23](#)
- [Quality of Service, page 2-25](#)

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

## Configuring 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.

## 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.

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

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

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

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | From Fabric Manager, choose <b>ISLs &gt; FCIP</b> from the Physical Attributes pane. |
|---------------|--|

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

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 Fabric Manager 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.



### 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.

---

## 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.

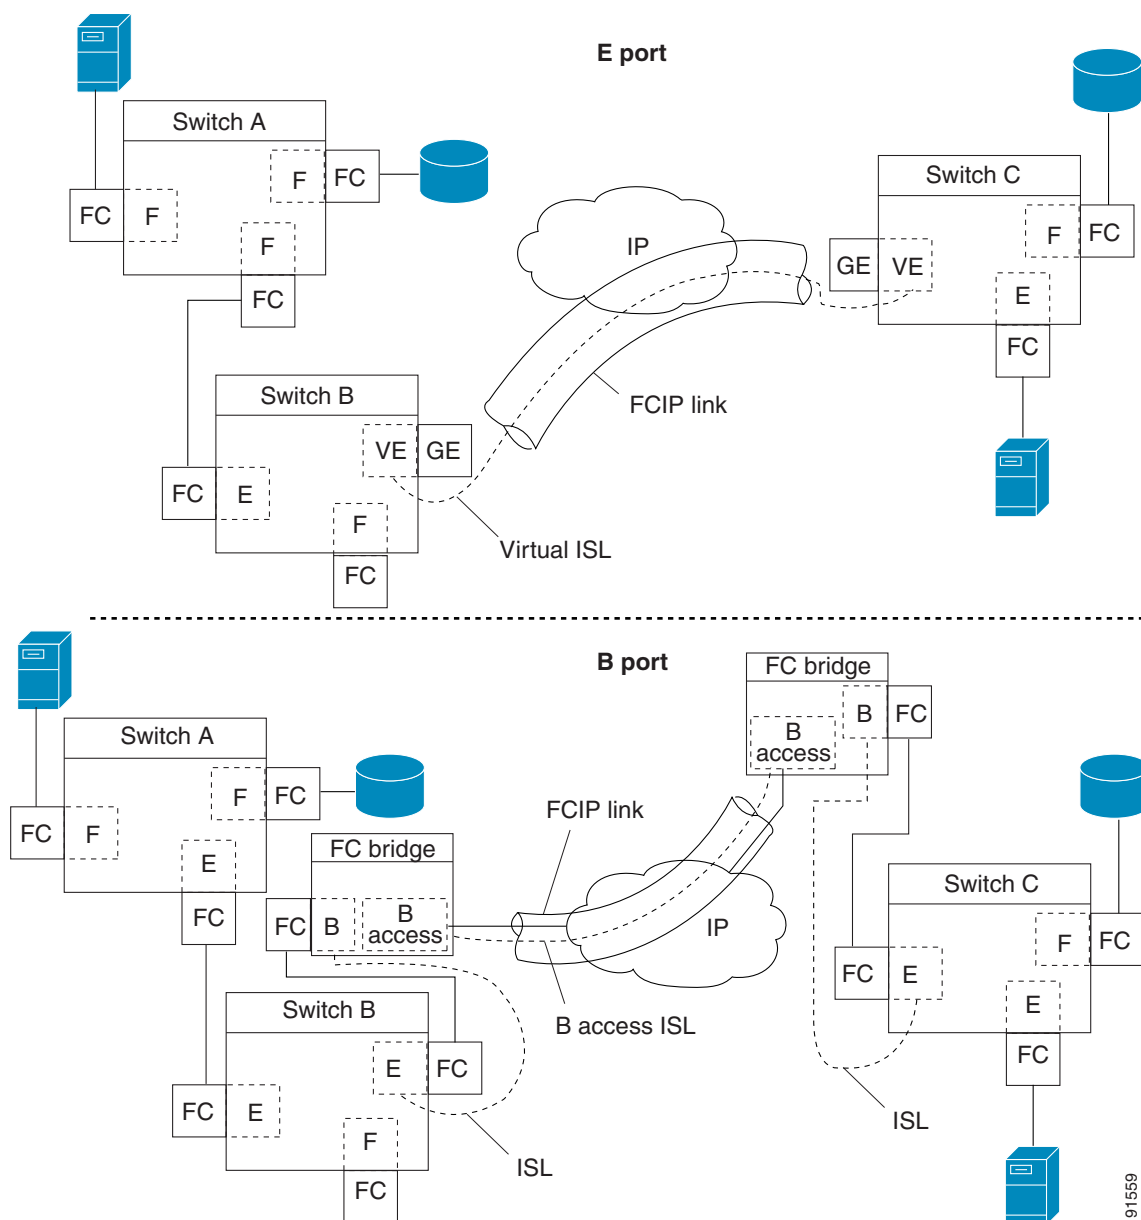
---

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

## 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 2-20](#) shows a typical SAN extension over an IP network.

**Figure 2-20 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

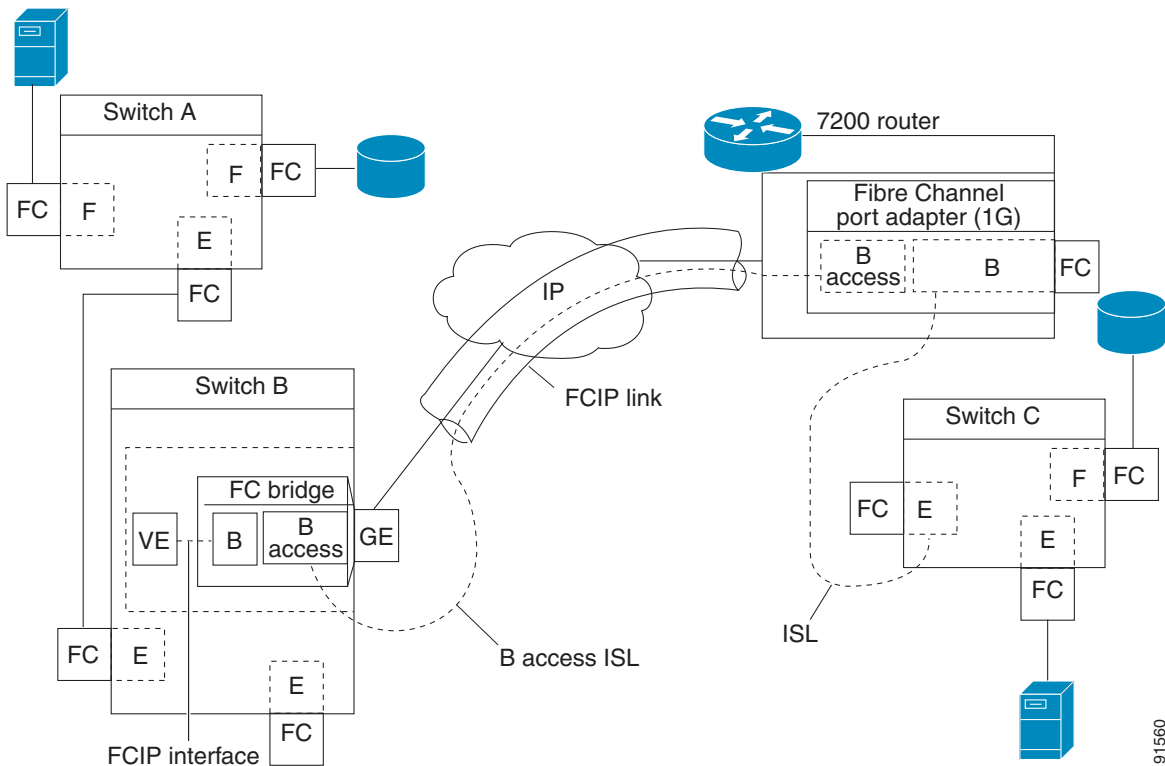
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

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*.

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 2-21](#)).

**Figure 2-21 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.

## 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

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

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.

To enable B port mode using Fabric Manager, follow these steps:

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

## 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.

## Configuring E Ports

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

- An FCIP interface can be a member of any VSAN
- Trunk mode and trunk allowed VSANs
- PortChannels
- FSPF

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

- 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 *Cisco Fabric Manager Fabric Configuration Guide*.
- Trunk mode and trunk allowed VSANs  
See the *Cisco Fabric Manager 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 *Cisco Fabric Manager Security Configuration Guide*.
- FSPF  
See the *Cisco Fabric Manager Fabric Configuration Guide*
- Fibre Channel domains (fcdomains)  
See the *Cisco Fabric Manager System Management Configuration Guide*.
- Importing and exporting the zone database from the adjacent switch  
See the *Cisco Fabric Manager System Management Configuration Guide*.

## Advanced FCIP Features

You can significantly improve application performance by configuring one or more of the following options for the FCIP interface:

- [FCIP Write Acceleration, page 2-26](#)
- [Configuring FCIP Write Acceleration, page 2-28](#)
- [FCIP Tape Acceleration, page 2-28](#)
- [Configuring FCIP Tape Acceleration, page 2-33](#)
- [FCIP Compression, page 2-33](#)

## 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.



### 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.

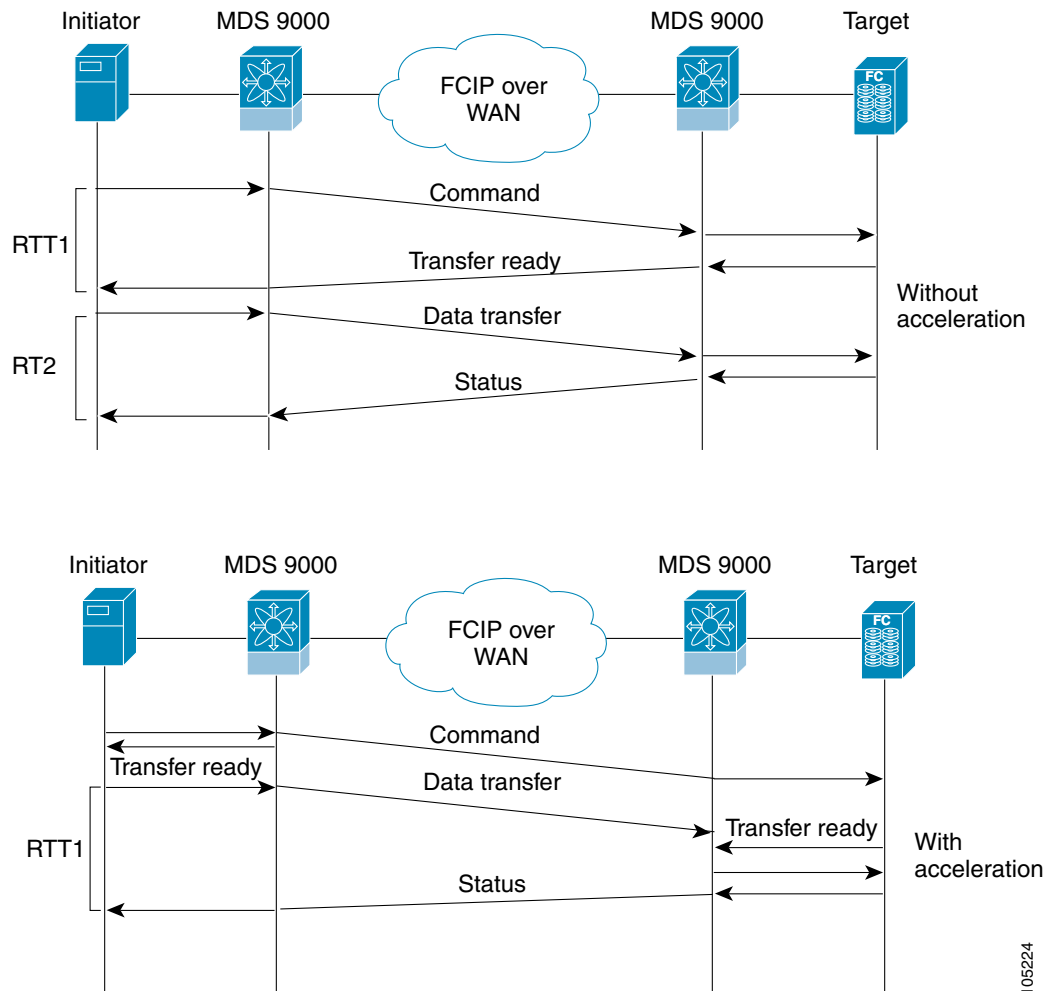
The WRITE command (see [Figure 2-22](#)), without write acceleration requires two round-trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE command



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

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 2-22 FCIP Link Write Acceleration**



**Tip**

FCIP write acceleration can be enabled for multiple FCIP tunnels if the tunnels are part of a dynamic PortChannel configured with channel mode active. FCIP write acceleration does not work if multiple non-PortChannel ISLs exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or failed WRITE or READ operations.



**Tip**

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

105224

*Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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 Port Channels. Also, FCIP write acceleration can be used in Port Channels configured with channel mode active or constructed with Port Channel Protocol (PCP).



**Caution**

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.

### Configuring FCIP Write Acceleration

You can enable FCIP write acceleration when you create the FCIP link using the FCIP Wizard. To enable write acceleration on an existing FCIP link, follow these steps:

- Step 1

Choose **ISLs > FCIP** from the Physical Attributes pane on Fabric Manager.

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 FICP link information (see [Figure 2-23](#)).

**Figure 2-23** *FCIP Tunnels (Advanced) Tab*

Switch	ProfileId	Interface	Timestamp Enable	Timestamp Tolerance	NumConn	Passive	QoS Control	QoS Data	IP Compression	Write Accelerator	Write Accelerator Oper
sw172-22-46-174	3	fcip3	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	none	<input checked="" type="checkbox"/>	false
sw172-22-46-174	4	fcip4	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	none	<input type="checkbox"/>	false
sw172-22-46-174	7	fcip7	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	high-comp-ratio(1.3)	<input type="checkbox"/>	false
sw172-22-46-174	8	fcip9	<input type="checkbox"/>	2000	2	<input type="checkbox"/>	0	0	high-throughput(1.3)	<input type="checkbox"/>	false

- 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.

### FCIP Tape Acceleration

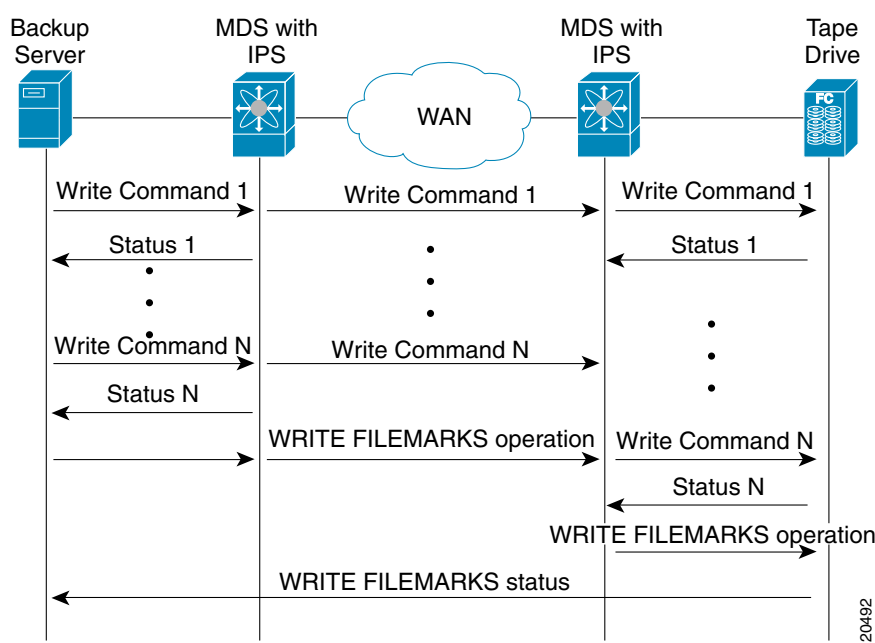
The FCIP write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.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

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

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 2-24](#)) 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.

**Figure 2-24 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**

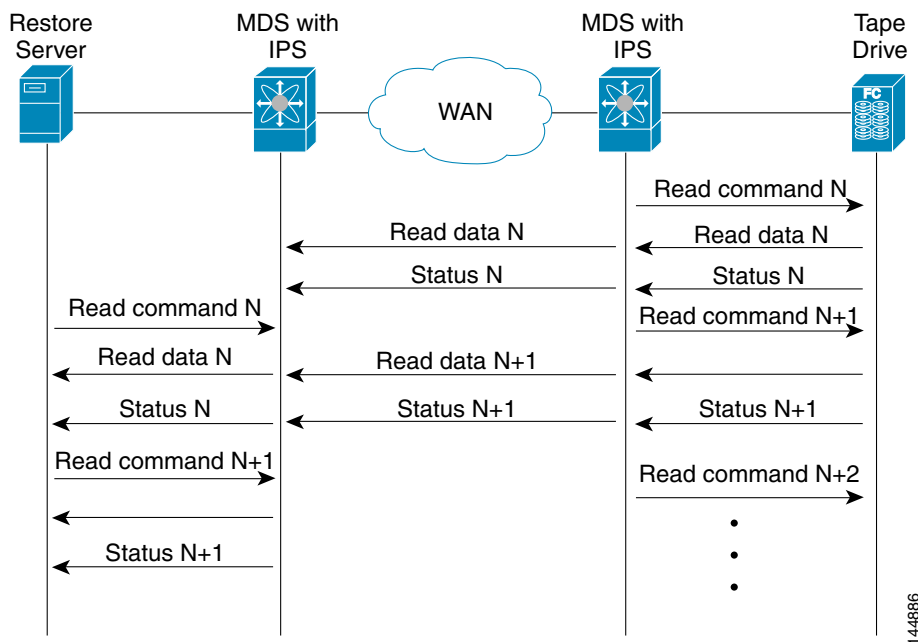
In some cases such as a quick link up/down event (FCIP link, Server/Tape Port link) in a tape library environment that exports Control LUN or a Medium Changer as LUN 0 and tape drives as other LUNs, tape acceleration may not detect the tape sessions and may not accelerate these sessions. You need to 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.

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

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 2-25) 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.

**Figure 2-25 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.

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

**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 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.

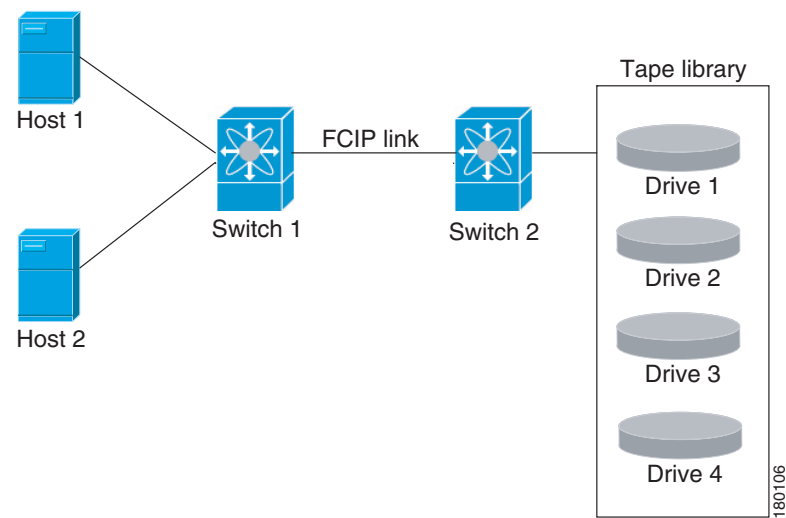
### Tape Library LUN Mapping for FCIP Tape Acceleration

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 2-26](#) 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.

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

**Figure 2-26      FCIP LUN Mapping Example**



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](#).

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

**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

## Configuring FCIP Tape Acceleration

To enable FCIP tape acceleration using Fabric Manager, follow these steps:

- 
- Step 1** From Fabric Manager, 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 Fabric Manager 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.
- 

## 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).

Mode1 and Mode3 compression modes are deprecated in Cisco MDS NX-OS Release 5.0(1a) and later. The MSM-18/4 and SSN-16 modules support Auto and Mode2 compression modes. Both of these modes internally use the hardware compression engine in the module. Auto mode is enabled by default. Mode2 uses a larger batch size for compression than Auto-mode, which results in higher compression throughput. However, Mode2 incurs a small latency due to the compression throughput. For those deployments where aggressive throughput is most important, Mode2 can be used.



### 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***

If both ends of the FCIP link are running NX-OS Release 4.x and 5.0(1a) and later, 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 2-4 lists the default settings for FCIP parameters.

**Table 2-4 Default FCIP Parameters**

Parameters	Default
TCP default port for FCIP	3225
<b>minimum-retransmit-time</b>	200 msec
Keepalive timeout	60 sec
Maximum retransmissions	4 retransmissions
PMTU discovery	Enabled
<b>pmtu-enable reset-timeout</b>	3600 sec
SACK	Enabled
<b>max-bandwidth</b>	1 Gbps
<b>min-available-bandwidth</b>	500 Mbps
<b>round-trip-time</b>	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
<b>special-frame</b>	Disabled
FCIP timestamp	Disabled
<b>acceptable-diff</b> range to accept packets	+/- 2000 msec
B port keepalive responses	Disabled
Write acceleration	Disabled
Tape acceleration	Disabled





## CHAPTER 3

# Configuring the SAN Extension Tuner

---

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

This chapter includes the following sections:

- [About the SAN Extension Tuner, page 3-1](#)
- [License Prerequisites, page 3-3](#)
- [Configuring the SAN Extension Tuner, page 3-4](#)
- [Using the SAN Extension Tuner Wizard, page 3-4](#)
- [Default Settings, page 3-6](#)

## 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.

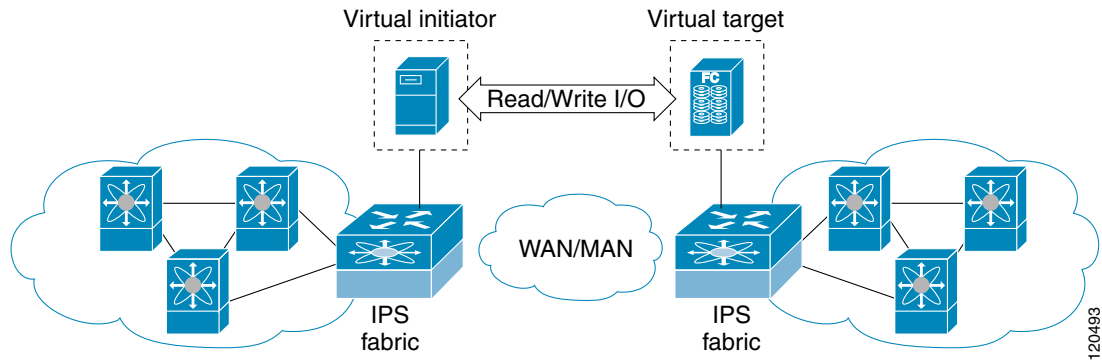
Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. To achieve maximum throughput performance across the fabric, you can tune the following configuration parameters:

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

- The TCP parameters for the FCIP profile (see “Window Management” section on page 2-19 for more information).
- 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 4-6 for more information.
- Configure the virtual N ports in a separate VSAN or zone as required by your network.
  - Be aware that a separate VSAN with only virtual N ports is not required, but is recommended as some legacy HBAs may fail if logins to targets are rejected.
  - Do not use same Gigabit Ethernet interface to configure virtual N ports and FCIP links—use different Gigabit Ethernet interfaces. While this is not a requirement, it is recommended as the traffic generated by the virtual N ports may interfere with the performance of the FCIP link.

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

## 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**



## Data Pattern

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

## License Prerequisites

To use the SET, you need to obtain the SAN\_EXTN\_OVER\_IP license (see the *Cisco Family NX-OS Licensing Guide*).

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

# Configuring the SAN Extension Tuner

This section includes the following topics:

- [Tuning the FCIP Link, page 3-4](#)

## Tuning the FCIP Link

To tune the required FCIP link, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Configure the nWWN for the virtual N ports on the switch.   |
| <b>Step 2</b> | Enable iSCSI on the interfaces on which you want to create the N ports.   |
| <b>Step 3</b> | Configure the virtual N ports on either side of the FCIP link.  |
| <b>Step 4</b> | Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see the <i>Cisco Fabric Manager 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. |
| <b>Step 5</b> | Start the SCSI read and write I/Os.   |
| <b>Step 6</b> | 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

Use the SAN Extension Tuner wizard to perform the these tasks:

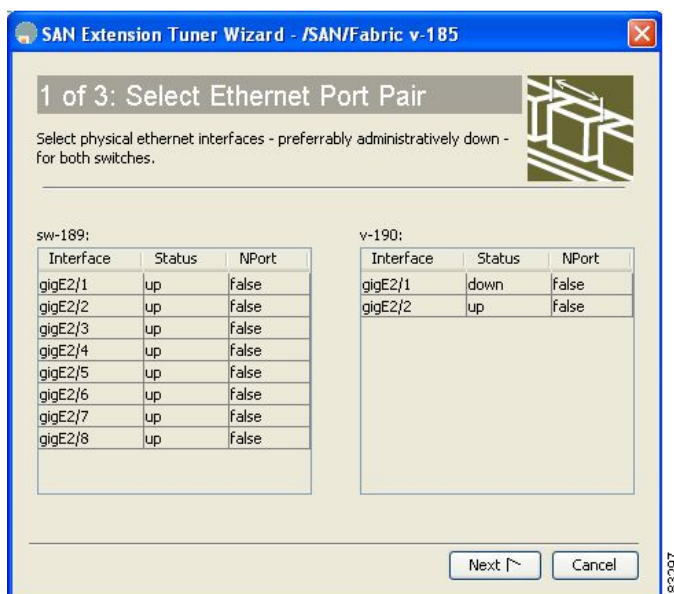
- Configuring nWWN ports
- Enabling iSCSI
- Configuring Virtual N ports
- Assigning SCSI read and write CLI commands
- Assigning SCSI tape read and write CLI commands
- Configuring a data pattern for SCSI commands

To tune the required FCIP link using the SAN Extension Tuner Wizard in Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Right-click a valid FCIP link in the Fabric pane, and then select <b>SAN Extension Tuner</b> from the drop-down list. You can also highlight the link and choose <b>Tools &gt; Other &gt; SAN Extension Tuner</b> .<br>You see the Select Ethernet Port Pair dialog box (see <a href="#">Figure 3-4</a> ). |
|---------------|--|

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

**Figure 3-4 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**.

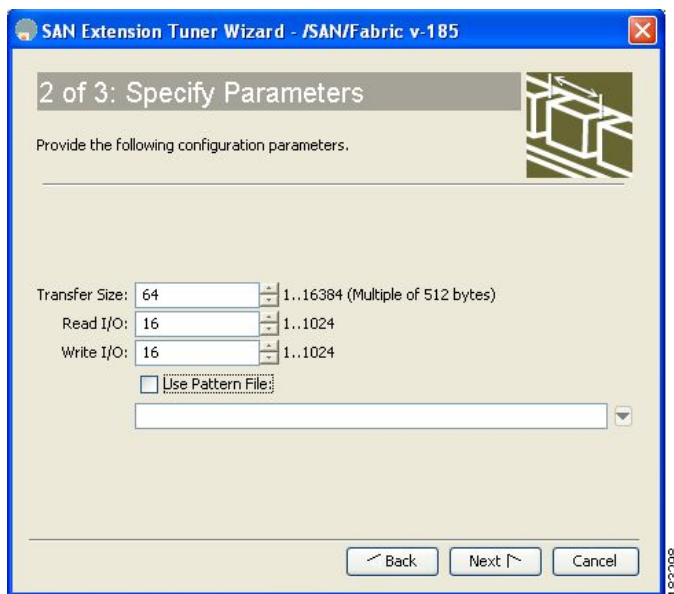


**Note** The Ethernet ports you select should be listed as down.

You see the Specify Parameters dialog box (see [Figure 3-5](#)).

- 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.

**Figure 3-5 Specify Parameters Dialog Box**



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

- Step 4** (Optional) Change the default settings for the transfer data size and the number of concurrent SCSI read and write commands as follows:
- Set Transfer Size to the number of bytes that you expect your applications to use over the FCIP link.
  - Set Read I/O to the number of concurrent SCSI read commands you expect your applications to generate over the FCIP link.
  - Set Write I/O to the number of concurrent outstanding SCSI write commands you expect your applications to generate over the FCIP link.



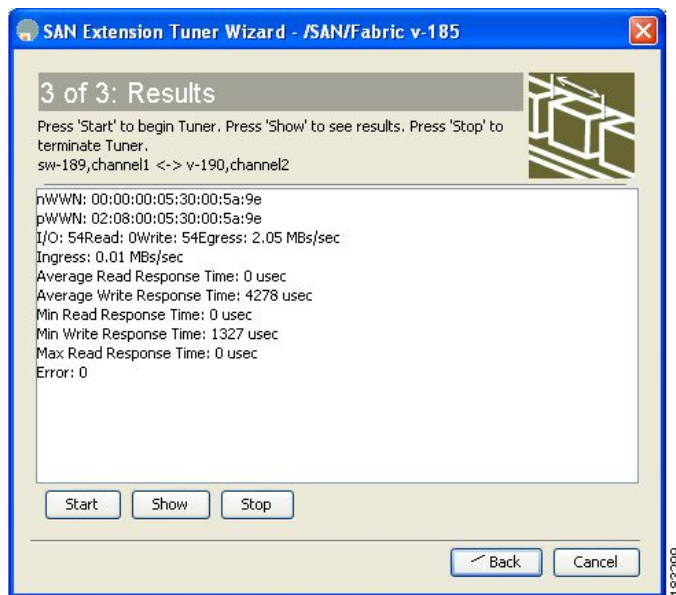
**Note** There is only one outstanding I/O at a time to the virtual N port that emulates the tape behavior.

- 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 3-3](#).

- Step 5** Click **Next**.

You see the Results dialog box (see [Figure 3-6](#)).

**Figure 3-6 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.

## Default Settings

[Table 3-1](#) lists the default settings for tuning parameters.

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

**Table 3-1**      **Default Tuning Parameters**

Parameters	Default
Tuning	Disabled
Transfer ready size	Same as the transfer size in the SCSI <b>write</b> command
Outstanding I/Os	1
Number of transactions	1
Data generation format	All-zero format
File mark frequency	0

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





## CHAPTER 5

# Configuring IP Services

---

## 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.



### Note

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

This chapter includes the following sections:

- [Traffic Management Services, page 5-2](#)
- [Management Interface Configuration, page 5-2](#)
- [Default Gateway, page 5-3](#)
- [IPv4 Default Network Configuration, page 5-6](#)
- [IPFC, page 5-7](#)
- [IPv4 Static Routes, page 5-7](#)
- [Overlay VSANs, page 5-8](#)
- [Configuring Multiple VSANs, page 5-9](#)
- [Virtual Router Redundancy Protocol, page 5-10](#)
- [DNS Server Configuration, page 5-14](#)

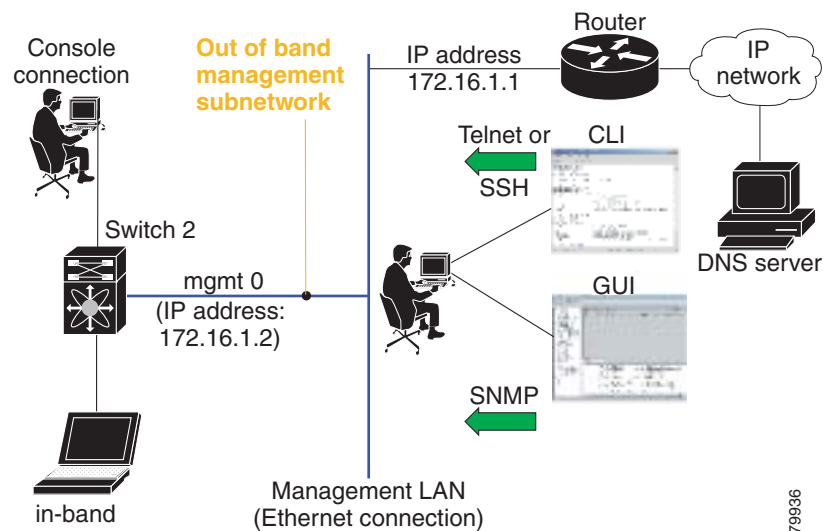
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

- [Default Settings, page 5-15](#)

## 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**



## 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

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

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.

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.
- 

## Default Gateway

You can configure a default gateway IPv4 address on your Cisco MDS 9000 Family switch.

This section includes the following topics:

- [About the Default Gateway, page 5-3](#)
- [Configuring the Default Gateway, page 5-3](#)

## About the Default Gateway

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 (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.

## Configuring the Default Gateway

To configure an IP route using Fabric Manager, follow these steps:

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

**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 as shown in [Figure 5-2](#).

**Figure 5-2 IP Route For Multiple Switches**

Switch	Destination, Mask, Gateway	Metric	Interface	Active
sw172-22-46-221	default, 0, 172.22.46.1	1	mgmt0	true
sw172-22-46-182	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-46-224	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-47-167	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-47-132	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-46-222	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-46-225	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-46-223	default, 0, 172.22.46.1	1	mgmt0	true
sw172-22-46-174	default, 0, 172.22.46.1	1	mgmt0	true
sw172-22-47-133	default, 0, 172.22.46.1	0	mgmt0	true
sw172-22-46-233	default, 0, 172.22.46.1	0	mgmt0	true

**Step 3** Click the **Create Row** icon to add a new IP route.

You see the dialog box shown in [Figure 5-3](#).

**Figure 5-3 User-Defined Command Dialog Box**

**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:

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


**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**.

You see the dialog box shown in [Figure 5-4](#).

**Figure 5-4** User-Defined Command Dialog Box



**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.




**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.

The new IP route is created as shown in [Figure 5-5](#).

**Figure 5-5** IP Route Window



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



#### 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.
```

## 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 5-6](#).

**Figure 5-6 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.

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

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.

## 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.

## IPFC Configuration Guidelines

Follow these guidelines to configure IPFC:

1. Create the VSAN to use for in-band management, if necessary.
2. Configure an IPv4 address and subnet mask for the VSAN interface.
3. Enable IPv4 routing.
4. Verify connectivity.

## IPv4 Static Routes

If your network configuration does not need an external router, you can configure IPv4 static routing on your MDS switch.



### Note

For information about IPv6 static routing, see the [Chapter 8, “Configuring IPv6 for Gigabit Ethernet Interfaces.”](#)

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.

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

## Overlay VSANs

This section describes overlay VSANs and how to configure them.

This section includes the following topics:

- [About Overlay VSANs, page 5-8](#)
- [Configuring Overlay VSANs, page 5-8](#)

## About Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

## Configuring Overlay VSANs

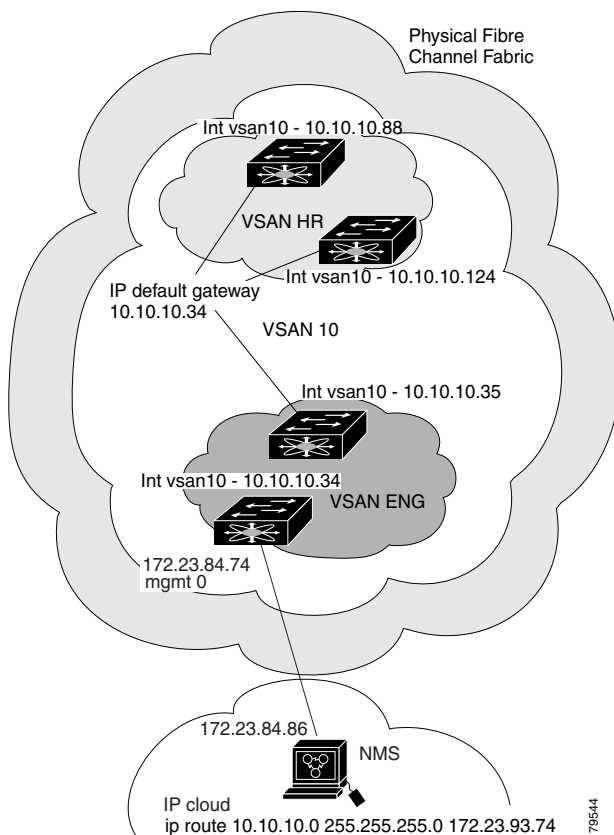
To configure an overlay VSAN, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Add the VSAN to the VSAN database on all switches in the fabric.  |
| <b>Step 2</b> | 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. |
| <b>Step 3</b> | Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.  |
| <b>Step 4</b> | Configure the default gateway (route) and the IPv4 address on switches that point to the NMS as shown in <a href="#">Figure 5-7</a> .   |



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

**Figure 5-7 Overlay VSAN Configuration Example**



## Configuring Multiple VSANs

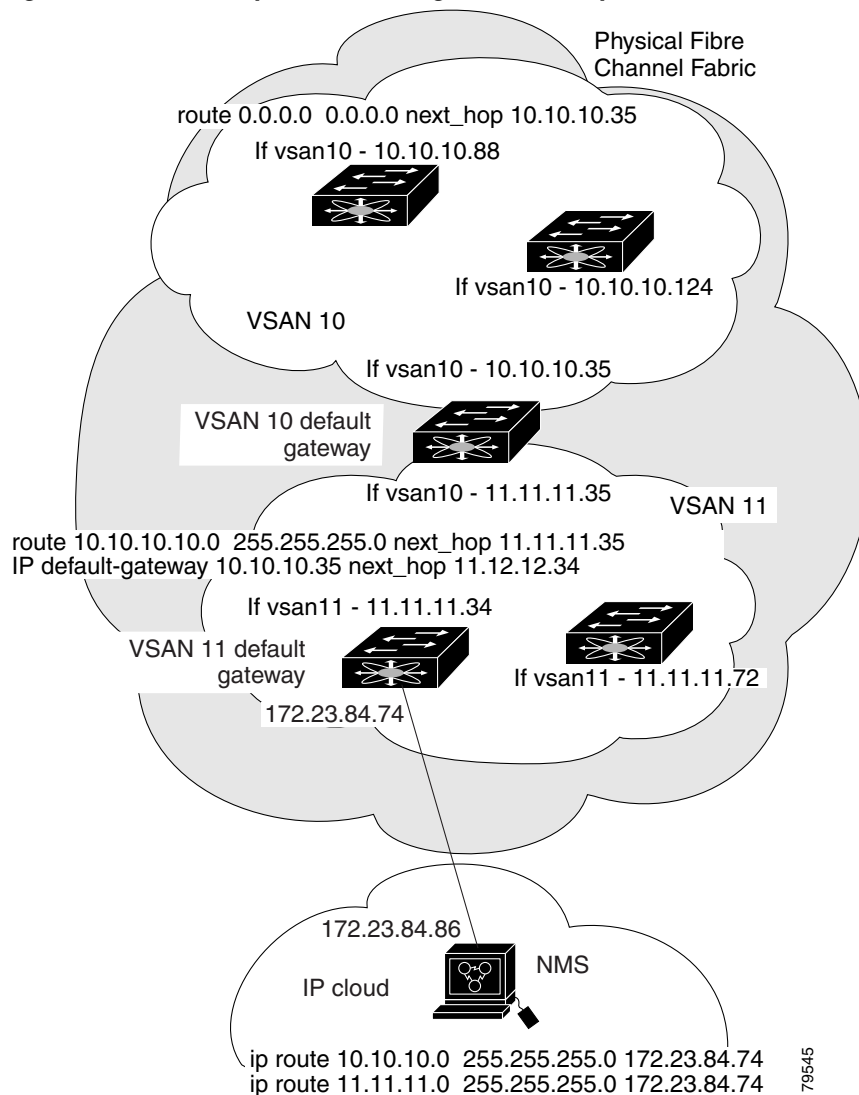
More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure 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 5-8](#).

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

**Figure 5-8 Multiple VSAN Configuration Example**



## Virtual Router Redundancy Protocol

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. This section provides details on the VRRP feature.

This section includes the following topics:

- [About VRRP, page 5-11](#)
- [Configuring VRRP, page 5-12](#)

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

## 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 characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338 and the draft-ietf-vrrp-ipv6 specification.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- 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 5-9](#), 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.

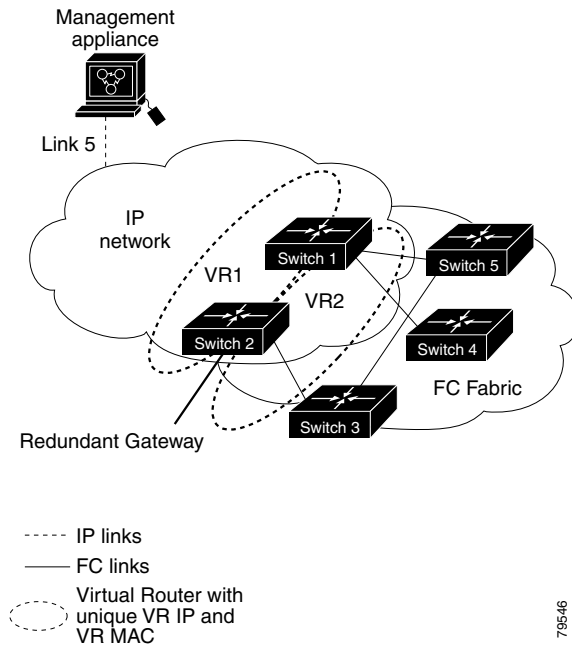
**Figure 5-9 VRRP Functionality**



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

In [Figure 5-10](#), the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

**Figure 5-10 Redundant Gateway**



## Configuring VRRP

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

- [Adding and Deleting a Virtual Router, page 5-12](#)
- [Virtual Router Initiation, page 5-13](#)
- [Adding Virtual Router IP Addresses, page 5-13](#)
- [Setting the Priority for the Virtual Router, page 5-13](#)
- [Setting the Time Interval for Advertisement Packets, page 5-13](#)
- [Configuring or Enabling Priority Preemption, page 5-13](#)
- [Setting Virtual Router Authentication, page 5-14](#)
- [Tracking the Interface Priority, page 5-14](#)

## 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.

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

## 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.

## Adding Virtual Router IP Addresses

One virtual router IP address can be configured for a virtual router. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address. You can configure either an IPv4 address or an IPv6 address.

According to the VRRP specification, the master VRRP router drops the packets addressed to the virtual router's IP address because the virtual router is only intended as a next-hop router to forward packets. In MDS switches however, some applications require that packets addressed to virtual router's IP address be accepted and delivered to them. By using the **secondary** option to the virtual router IPv4 address, the VRRP router will accept these packets when it is the master.

To 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**.
- 

## 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.

## 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.

## 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.

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

## 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.

## 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 5-13](#)). When the tracked interface is up, the priority of the virtual router is restored to the interface state tracking value. You can track the state of either a specified VSAN interface or the management interface (mgmt 0). The interface state tracking feature is disabled by default.



**Note**

For interface state tracking to function, you must enable preemption on the interface. See the [“Configuring or Enabling Priority Preemption” section on page 5-13](#).

## 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.

***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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

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





## CHAPTER 4

# 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 [“Basic Gigabit Ethernet Configuration for IPv4” section on page 7-2](#).

This chapter includes the following sections:

- [About iSCSI, page 4-1](#)
- [Configuring iSCSI, page 4-4](#)
- [Configuring iSLB, page 4-36](#)
- [iSCSI High Availability, page 4-51](#)
- [iSCSI Authentication Setup Guidelines and Scenarios, page 4-57](#)
- [iSNS, page 4-68](#)
- [iSNS Cloud Discovery, page 4-74](#)
- [Default Settings, page 4-76](#)

## 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

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

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.

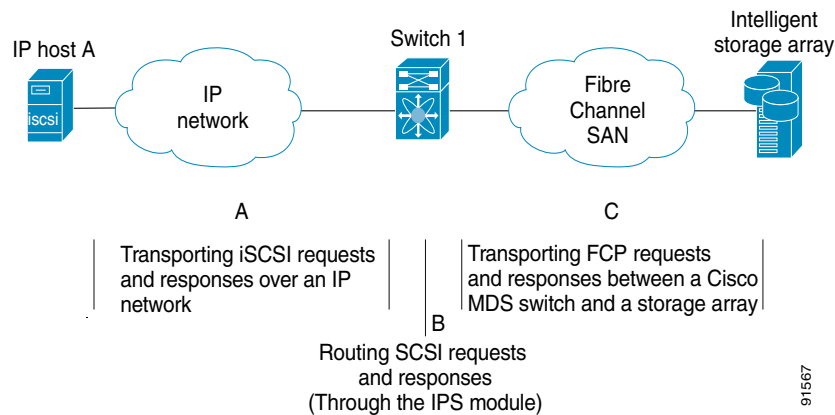


**Note**

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. (The Cisco.com website at <http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420-scsi> provides a list of compatible drivers). Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be 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](#)).

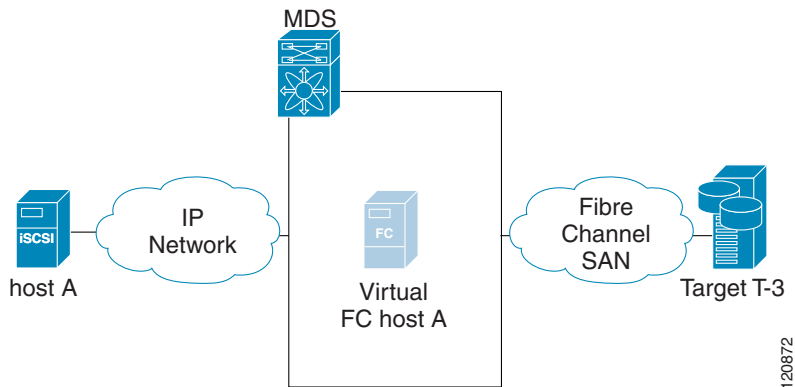
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**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 [fm-docfeedback@cisco.com](mailto:fm-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.

## 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 4-5](#)
- [Creating iSCSI Interfaces, page 4-6](#)
- [Using the iSCSI Wizard, page 4-6](#)
- [Presenting Fibre Channel Targets as iSCSI Targets, page 4-8](#)
- [Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 4-15](#)
- [iSCSI Access Control, page 4-25](#)
- [iSCSI Session Authentication, page 4-29](#)
- [iSCSI Immediate Data and Unsolicited Data Features, page 4-32](#)
- [iSCSI Interface Advanced Features, page 4-33](#)
- [Displaying iSCSI Information, page 4-46](#)

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

## 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 Fabric Manager 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.

To enable iSCSI on any switch using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).

**Figure 4-5** *iSCSI Tables in Fabric Manager*

Switch	Feature	Status	Command	LastCommand	Result
sw172-22-46-182	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-222	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-224	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-233	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-220	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-182	iscsi-interface-vsan-membership	disabled	noSelection	noSelection	none
sw172-22-46-221	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-224	iscsi-interface-vsan-membership	disabled	noSelection	noSelection	none
sw172-22-46-223	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-225	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-174	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-153	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-153	iscsi-interface-vsan-membership	enabled	noSelection	noSelection	none

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

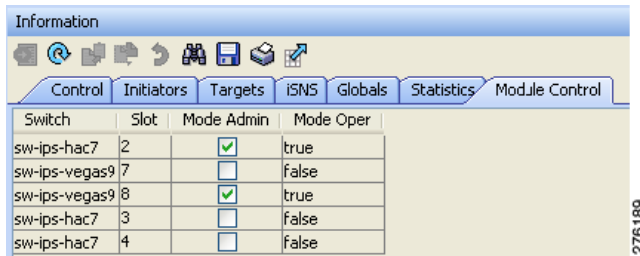
- Step 2** Choose **enable** from the Command column for each switch that you want to enable iSCSI on.  
**Step 3** Click the **Apply Changes** icon to save these changes.

To enable iSCSI on a module using Fabric Manager, 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 **Module Control** tab.  
You see the Module Control dialog box in the information pane (see [Figure 4-6](#)).

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

**Figure 4-6** *Module Control Dialog Box*

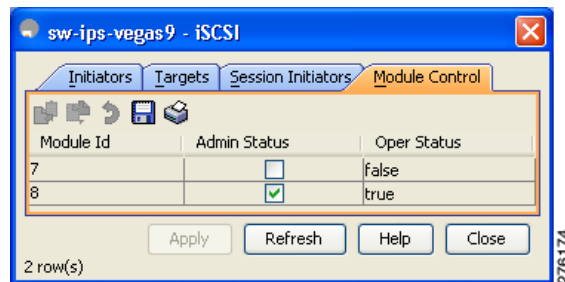


- Step 3** Check the **Mode Admin** check box to enable iSCSI for a specified port on the selected module.
- Step 4** Click the **Apply Changes** icon to save these changes.

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

- Step 1** Choose **IP > iSCSI**
- You see the iSCSI table (see [Figure 4-7](#)).

**Figure 4-7** *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.

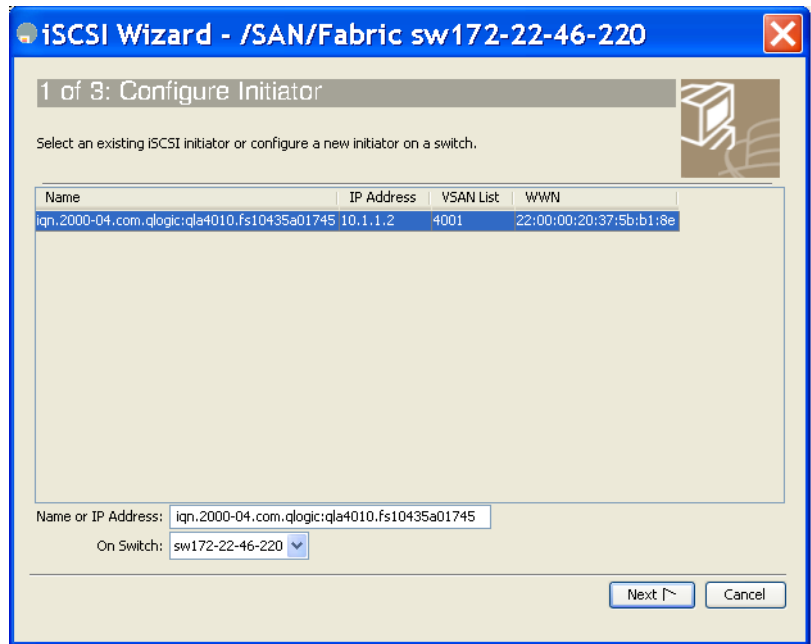
## Using the iSCSI Wizard

To use the iSCSI wizard in Fabric Manager, follow these steps:

- Step 1** Click the **iSCSI Setup Wizard** icon.
- You see the iSCSI Wizard Configure Initiator dialog box (see [Figure 4-8](#)).

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

**Figure 4-8 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 (see [Figure 4-9](#)).

**Figure 4-9 iSCSI Wizard Select Targets Dialog Box**



**Step 4** Select the VSAN and targets to associate with this iSCSI initiator and click **Next**.

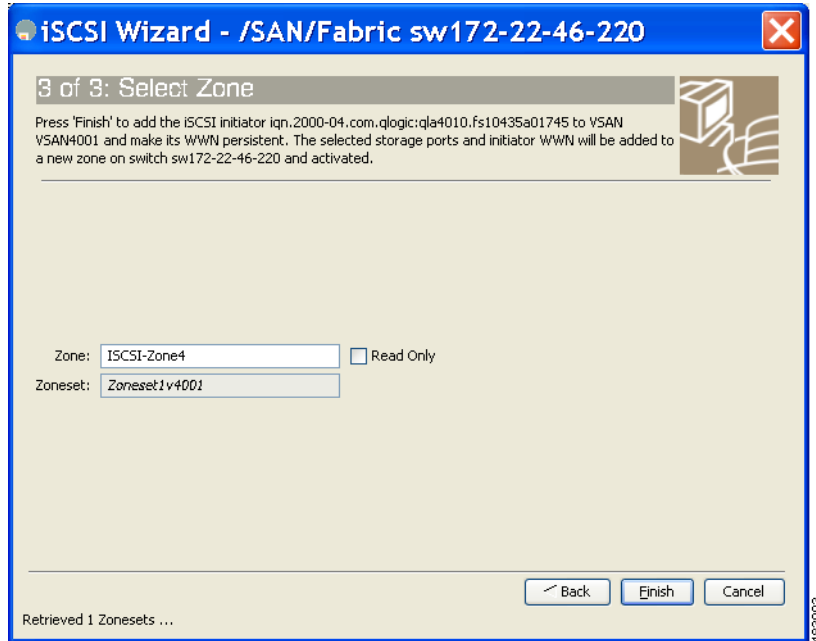
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Note**

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

You see the iSCSI Wizard Select Zone dialog box (see [Figure 4-10](#)).

**Figure 4-10 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.

## 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.



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

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 4-25](#)). 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 4-51](#)).



**Note**

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

## Dynamic Mapping

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

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

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

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

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

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

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

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

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



**Note**


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

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

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

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

Figure 4-11 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 4-25).

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

Step 1 Choose IP > iSCSI.

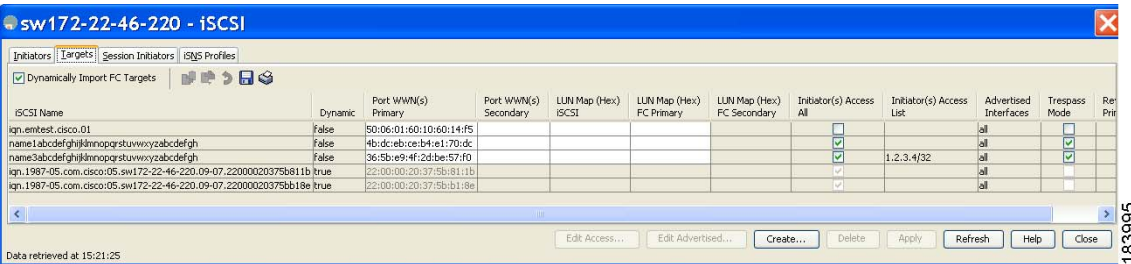
You see the iSCSI configuration (see Figure 4-12).

Figure 4-12 iSCSI Configuration in Device Manager



Step 2 Click the Target tab to display a list of existing iSCSI targets (see Figure 4-13).

Figure 4-13 iSCSI Targets Tab



Step 3 Check the Dynamically Import FC Targets check box.

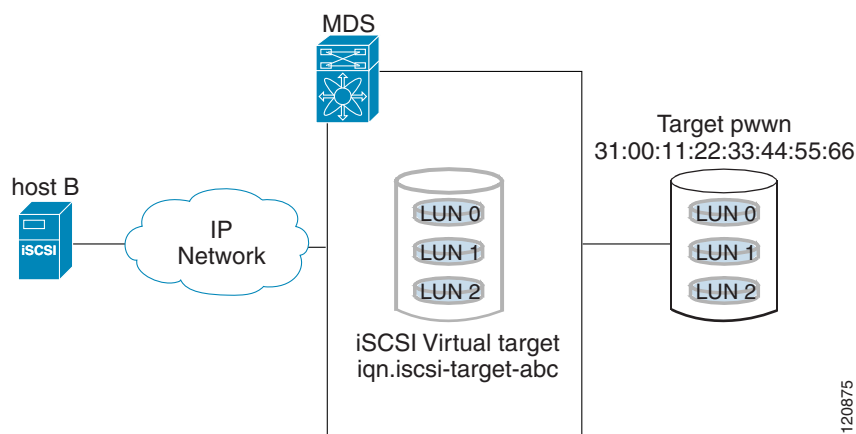
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Step 4** Click **Apply** to save this change.

## 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 4-14](#)).

**Figure 4-14**      **Statically Mapped iSCSI Targets**



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 (see [Figure 4-12](#)).

**Step 2** Click the **Targets** tab to display a list of existing iSCSI targets (see [Figure 4-13](#)).

**Step 3** Click **Create** to create an iSCSI target.

You see the Create iSCSI Targets dialog box (See [Figure 4-15](#)).

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

**Figure 4-15 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 4-25.
- 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 4-27 for more information on controlling access to statically mapped targets.

***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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.

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 (see [Figure 4-12](#)).

**Step 2** Click the **Targets** tab to display a list of existing iSCSI targets (see [Figure 4-13](#)).

**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.

---


## iSCSI Virtual Target Configuration Examples

This section provides three examples of iSCSI virtual target configurations.

### Example 1

This example assigns the whole Fibre Channel target as an iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 4-16](#)).

**Figure 4-16** Assigning iSCSI Node Names




### 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 4-17](#)).

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

**Figure 4-17 Mapping LUNs to an iSCSI Node Name**




```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```

### Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see [Figure 4-18](#)).

**Figure 4-18 Mapping LUNs to Multiple iSCSI Node Names**



```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
```

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

## Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

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

### Initiator Identification

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

- iSCSI qualified name (IQN)

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

- IP address

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

You can configure the iSCSI initiator identification mode on each IPS port and all the iSCSI hosts terminating on the IPS port will be identified according to that configuration. The default mode is to identify the initiator by name.

To specify the initiator identification mode using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Interfaces &gt; FC Logical</b> from the Physical Attributes pane.<br>You see the interfaces configuration in the Information pane.                |
| <b>Step 2</b> | Click the <b>iSCSI</b> tab.<br>You see the iSCSI interfaces configuration.  |
| <b>Step 3</b> | Right-click the Initiator ID Mode field for the iSCSI interface that you want to modify and select <b>name</b> or <b>ipaddress</b> from the drop-down menu. |
| <b>Step 4</b> | Click <b>Apply Changes</b> to save this change.   |
- 

## 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.

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

- 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 4-46.

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 4-19](#)). Every Fibre Channel N port requires a unique Node WWN and Port WWN.



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

**Figure 4-19 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 4-19), and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

## iSCSI Initiator Idle Timeout

iSCSI initiator idle timeout specifies the time for which the virtual Fibre Channel N port is kept idle after the initiator logs out from its last iSCSI session. The default value for this timer is 300 seconds. This is useful to avoid N ports logging in to and logging off of the Fibre Channel SAN as transient failure occurs in the IP network. This helps reduce unnecessary RSCNs being generated in the Fibre Channel SAN.

To configure the initiator idle timeout using Fabric Manager, follow these steps:

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

- 
- Step 1** Choose **End Devices** > **iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).
- 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.
- 

## WWN Assignment for iSCSI Initiators

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

- Dynamic mapping (default)
- Static mapping

### Dynamic Mapping

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

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

The following are three dynamic initiator modes are supported:

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

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



#### Note

Configuring dynamic initiator modes is supported only through the CLI, not through Device Manager or Fabric Manager.

### 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.

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

- 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 *Cisco Fabric Manager Fabric Configuration Guide* for more information). You should not use any previously assigned WWNs.

To configure static mapping for an iSCSI initiator using Device Manager, follow these steps:

**Step 1** Select **IP > iSCSI**.

You see the iSCSI configuration (see [Figure 4-12](#)). The Initiators tab is the default.

**Step 2** Click **Create** to create an iSCSI initiator.

You see the Create iSCSI Initiators dialog box (see [Figure 4-20](#)).

**Figure 4-20** Create iSCSI Initiators Dialog Box

**Step 3** Set the iSCSI node name or IP address and VSAN membership.

**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.

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

- 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 4-29](#).
- Step 9** Click **Create** to create this iSCSI initiator.



### Note

If the system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is saved to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a write erase, you must manually delete the WWN configuration from the ASCII file. Failing to do so can cause duplicate WWN assignments if the ASCII configuration file is reapplied on the switch.

### Making the Dynamic iSCSI Initiator WWN Mapping Static

After a dynamic iSCSI initiator has already logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping so this initiator uses the same mapping the next time it logs in.

You can convert a dynamic iSCSI initiator to static iSCSI initiator and make its WWNs persistent (see the [“Dynamic Mapping” section on page 4-18](#)).



### 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 Fabric Manager. In Fabric Manager or Device Manager, you must delete and then recreate this initiator to have the pWWNs static.

### Checking for WWN Conflicts

WWNs assigned to static iSCSI initiators by the system can be inadvertently returned to the system when an upgrade fails or you downgrade the system software (manually booting up an older Cisco MDS SAN-OS release without using the **install all** command). In these instances, the system can later assign those WWNs to other iSCSI initiators (dynamic or static) and cause conflicts.

You can address this problem by checking for and removing any configured WWNs that belong to the system whenever such scenarios occur.

To permanently keep the automatically assigned nWWN mapping using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).
- 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.

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


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

## 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.

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 4-21](#)). 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 4-11) with LUN mapping and iSCSI access control (see the [“iSCSI Access Control”](#) section on page 4-25).

**Figure 4-21 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.

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



### 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 4-46.

To configure the proxy initiator using Fabric Manager, follow these steps:

- Step 1** Expand **Switches**, expand **Interfaces**, and then select **FC Logical** in the Physical Attributes pane. You see the Interface tables in the Information pane (see [Figure 4-22](#)).

**Figure 4-22 FC Logical Interface Tables**

Switch	Interface	Mode Admin	Mode Oper	Port VSAN	Dynamic VSAN	Description	Speed Admin	Speed Oper	Rate Mode	Status Service	Status Admin	Status Oper	FailureCause	Was Enabled	LastC
sw172-22-46-233	fcip2	auto	E		1 n/a		auto	1 Gb	shared	in	up	up	none	true	2007/10/11 18:39:58
sw172-22-46-221	channel1	E	TE		1 n/a	To sw172-22-46-220	auto	2 Gb	shared	in	up	up	none	false	2007/10/11 18:39:58
sw172-22-47-20	channel1	E	TE		1 n/a	To sw172-22-46-174	auto	10 Gb	shared	in	up	up	none	false	2007/10/11 18:39:58
sw172-22-47-133	channel1	E	TE		1 n/a	To sw172-22-47-132	auto	8 Gb	shared	in	up	up	none	false	2007/10/11 18:39:58
sw172-22-46-223	channel2	E	TE		1 n/a	To sw172-22-46-220	auto	1 Gb	shared	in	up	up	trunkNotFullyActive	false	2007/10/11 18:39:58
sw172-22-46-223	fcip6	auto	E		1 n/a		auto	1 Gb	shared	in	up	up	none	true	2007/10/11 18:39:58
sw172-22-46-223	channel1	E	TE		1 n/a	To sw172-22-46-220	auto	2 Gb	shared	in	up	up	trunkNotFullyActive	false	2007/10/11 18:39:58
sw172-22-47-132	channel1	E	TE		1 n/a	To sw172-22-47-133	auto	8 Gb	shared	in	up	up	trunkNotFullyActive	false	2007/10/11 18:39:58
sw172-22-46-220	channel4	E	TE		1 n/a	To sw172-22-46-221	auto	2 Gb	shared	in	up	up	trunkNotFullyActive	false	2007/10/11 18:39:58

- Step 2** In Device Manager, select **Interface > Ethernet** and **iSCSI**.

You see the Ethernet Interfaces and iSCSI dialog box (See in [Figure 4-23](#)).

**Figure 4-23 Ethernet Interfaces and iSCSI Dialog Box**

Interface	Description	Mtu	Oper	PhysAddress	Admin	Oper	LastChange	Connector Present	CDP	IscsiAuthMethod	iSNS ProfileName	Promiscuous Mode	Auto Negotiate	Beacon Mode
gigE8/1		2300	n/a	00:05:30:01:80:3e	up	down	2007/05/25-12:48:25	false	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE8/2		2300	1 Gb	00:05:30:01:80:3f	up	up	2007/05/24-01:17:48	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/1		1500	1 Gb	00:05:30:00:a1:9a	up	up	2007/06/07-08:18:59	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/2		1500	1 Gb	00:05:30:00:a1:9b	up	up	2007/06/07-08:18:59	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/3		2300	1 Gb	00:05:30:00:a1:9c	up	up	2007/06/07-08:18:59	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/4		1500	1 Gb	00:05:30:00:a1:9d	up	up	2007/06/07-08:18:59	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/5		2300	1 Gb	00:05:30:00:a1:9e	up	up	2007/05/16-15:03:58	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/6		2300	1 Gb	00:05:30:00:a1:9f	up	up	2007/05/16-15:03:58	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/7		1500	1 Gb	00:05:30:00:a1:a0	up	up	2007/05/16-15:03:58	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
gigE9/8		1500	1 Gb	00:05:30:00:a1:a1	up	up	2007/05/16-15:03:58	true	<input checked="" type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

- Step 3** Click the **iSCSI** tab in either FM or DM.

You see the iSCSI interface configuration table (see [Figure 4-24](#)).

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

**Figure 4-24 iSCSI Tab in Device Manager**

sw172-22-46-220 - Ethernet Interfaces and iSCSI

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

Close

10 rows

10 rows

Apply

Refresh

Help

**Step 4** Check the **Proxy Mode Enable** check box.

**Step 5** Click the **Apply Changes** icon in Fabric Manager 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 4-25](#)).

## 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).

## 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.

To assign VSAN membership for iSCSI hosts using Fabric Manager, follow these steps:

**Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.

You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).

**Step 2** Click the **Initiators** tab.

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

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 4-46.

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 (see [Figure 4-23](#)).

**Step 2** Click the **iSCSI** tab.

You see the iSCSI interface configuration table (see [Figure 4-24](#)).

**Step 3** Double-click the PortVSAN column and modify the default port VSAN.

**Step 4** Click **Apply** to save these changes.

## Example of VSAN Membership for iSCSI Devices


[Figure 4-25](#) 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) connect to iSCSI interface 1/1.



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

**Figure 4-25 VSAN Membership for iSCSI Interfaces**



Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so they will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

## Advanced VSAN Membership for iSCSI Hosts

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

## iSCSI Access Control

Two 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 4-26](#)

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

- [iSCSI-Based Access Control](#), page 4-27
- [Enforcing Access Control](#), page 4-28

## 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.

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 *Cisco Fabric Manager 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 4-16](#)), 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 4-27](#)).


To add an iSCSI initiator to the zone database using Fabric Manager, follow these steps:

### Step 1 Choose **Zone > Edit Local Full Zone Database**.

You see the Edit Local Zone Database dialog box (see [Figure 4-26](#)).

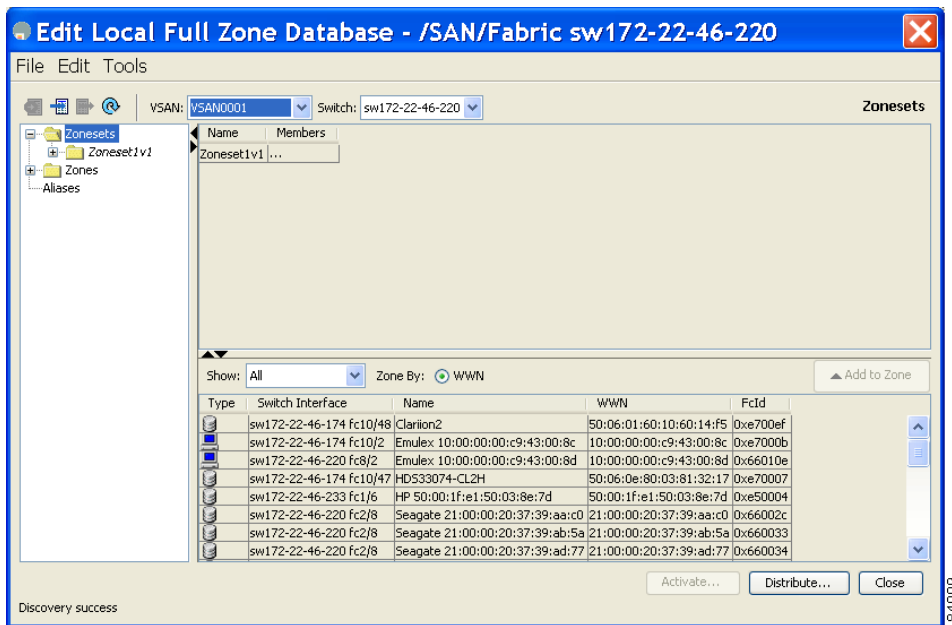
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 4-26** Edit Local Zone Database Dialog Box in Fabric Manager



- 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 (see [Figure 4-27](#)).

**Figure 4-27** Available Zones and Zone Sets



- 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.

## 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 4-11). 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 [fm-docfeedback@cisco.com](mailto:fm-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.

To configure access control in iSCSI using Device Manager, follow these steps:

- 
- Step 1** Select **IP > iSCSI**.  
You see the iSCSI configuration (see [Figure 4-12](#)).
- 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.
- Step 7** Click **Create** to add this initiator to the Initiator Access List.
- 

## 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 4-9).
- 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 4-27.

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.

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

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 *Cisco Fabric Manager Security Configuration Guide*.

To configure AAA authentication for an iSCSI user using Fabric Manager, 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 (see [Figure 4-28](#)).

**Figure 4-28 AAA per Application Configuration**

Switch	Type, SubType, Function	Server Group IdList	Local	Trivial
sw172-22-46-233	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-220	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-224	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-223	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-182	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-222	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-225	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-47-20	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-221	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-47-167	default, all, accounting		<input checked="" type="checkbox"/>	<input type="checkbox"/>
sw172-22-46-224	login, all, authentication		<input checked="" type="checkbox"/>	<input type="checkbox"/>

**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.

The following topics are included in this section:

- [Configuring Authentication Mechanism, page 4-30](#)

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

- [Configuring Local Authentication, page 4-31](#)
- [Restricting iSCSI Initiator Authentication, page 4-31](#)
- [Configuring Mutual CHAP Authentication, page 4-31](#)
- [Configuring an iSCSI RADIUS Server, page 4-32](#)

## 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.

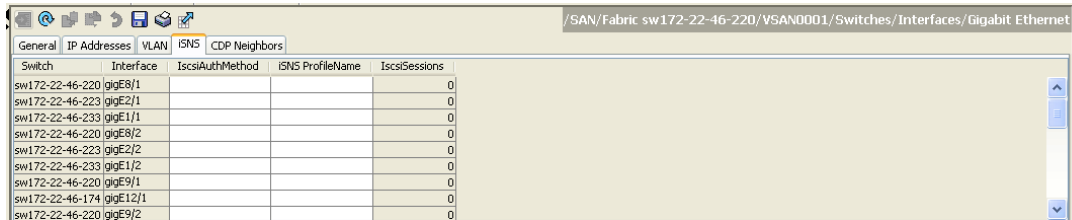
To configure AAA authentication for an iSCSI user using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).
- 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 Fabric Manager to save these changes.

To configure the authentication mechanism for iSCSI sessions to a particular interface using Fabric Manager, 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 (see [Figure 4-29](#)).

**Figure 4-29 Configuring iSCSI Authentication on an Interface**



Switch	Interface	IscsiAuthMethod	iSNS ProfileName	IscsiSessions
sw172-22-46-220	gigE8/1			0
sw172-22-46-223	gigE2/1			0
sw172-22-46-233	gigE1/1			0
sw172-22-46-220	gigE8/2			0
sw172-22-46-223	gigE2/2			0
sw172-22-46-233	gigE1/2			0
sw172-22-46-220	gigE9/1			0
sw172-22-46-174	gigE12/1			0
sw172-22-46-220	gigE9/2			0

- 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***


## Configuring Local Authentication

See the *Cisco Fabric Manager Security Configuration Guide* to create the local password database. To create users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory.

To configure iSCSI users for local authentication using Device Manager, follow these steps:

- 
- Step 1** Choose **Security > iSCSI**.  
You see the iSCSI Security dialog box (see [Figure 4-30](#)).

**Figure 4-30** 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.

To restrict an initiator to use a specific user name for CHAP authentication using Fabric Manager, follow these steps:

- 
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).
- 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.

## ***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)***

In addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator, the IPS module or MPS-14/2 module also supports a mechanism for the iSCSI initiator to authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication requires the user to configure a user name and password for the switch to present to the iSCSI initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

To configure a global iSCSI target user name and password to be used by the switch to authenticate itself to an initiator using Fabric Manager, follow these steps:

- 
- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.  
You see the iSCSI tables in the Information pane (see [Figure 4-5](#)).
  - Step 2** Select the **Globals** tab.  
You see the global iSCSI configuration.
  - Step 3** Fill in the Target UserName and Target Password fields.
  - 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 using Device Manager, follow these steps:

- 
- Step 1** Choose **IP > iSCSI**.  
You see the iSCSI configuration (see [Figure 4-12](#)).
  - 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**

To configure an iSCSI RADIUS server, follow these steps:

- 
- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
  - Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
  - Step 3** Configure the iSCSI users and passwords on the RADIUS server.
- 

## **iSCSI 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.



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

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

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

## iSCSI Interface Advanced Features

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. These configurations are similar to the advanced FCIP configurations and are already explained in that section.

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

- [iSCSI Listener Port, page 4-33](#)
- [TCP Tuning Parameters, page 4-33](#)
- [Setting QoS Values, page 4-34](#)
- [iSCSI Routing Modes, page 4-34](#)

### 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 [“Minimum Retransmit Timeout” section on page 2-18](#) for more information).
- Keepalive timeout (See the [“Keepalive Timeout” section on page 2-18](#) for more information).
- Maximum retransmissions (See the [“Maximum Retransmissions” section on page 2-18](#) for more information).
- Path MTU (See the [“Path MTUs” section on page 2-19](#) 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 [“Window Management” section on page 2-19](#) for more information).
- Buffer size (The iSCSI default send buffer size is 4096 KB) (See the [“Buffer Size” section on page 2-20](#) for more information).
- Window congestion monitoring (enabled by default and the default burst size is 50 KB) (See the [“Monitoring Congestion” section on page 2-19](#) for more information).
- Maximum delay jitter (enabled by default and the default time is 500 microseconds).

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

## Setting QoS Values

To set the QoS values using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Expand <b>Switches</b> , expand <b>Interfaces</b> and then select <b>FC Logical</b> in the Physical Attributes pane.<br>You see the Interface tables in the Information pane (see <a href="#">Figure 4-22</a> ). |
| <b>Step 2</b> | In Device Manager, choose <b>Interface &gt; Ethernet and iSCSI</b> .<br>You see the Ethernet Interfaces and iSCSI dialog box (see <a href="#">Figure 4-23</a> ).   |
| <b>Step 3</b> | Click the <b>iSCSI TCP</b> tab in either Fabric Manager or Device Manager.<br>You see the iSCSI TCP configuration table.   |
| <b>Step 4</b> | Set the QoS field from 1 to 6.   |
| <b>Step 5</b> | Click the <b>Apply Changes</b> icon in Fabric Manager or click <b>Apply</b> in Device Manager to save these changes.   |
- 

## 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.

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

- Cut-through mode

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

Figure 4-31 compares the messages exchanged by the iSCSI routing modes.

**Figure 4-31 iSCSI Routing Modes**

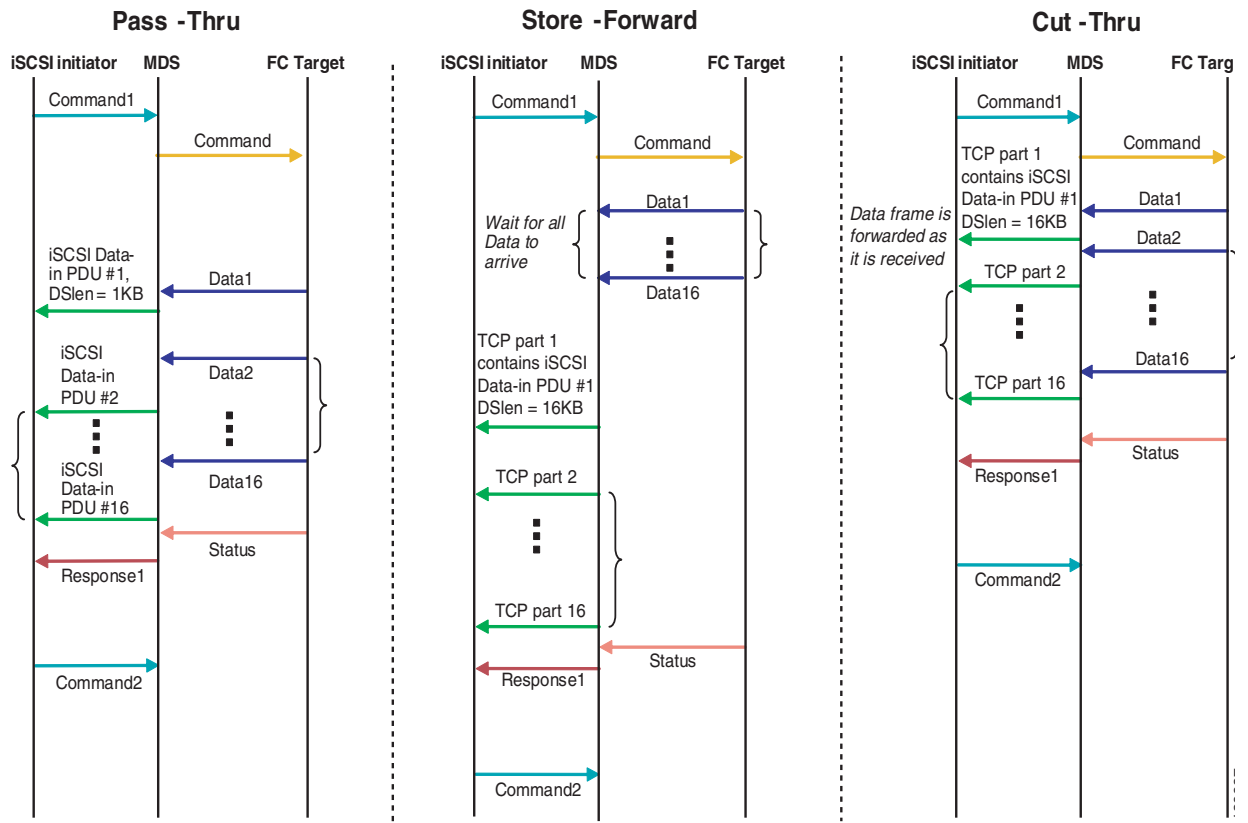


Table 4-1 compares the advantages and disadvantages of the different iSCSI routing modes.

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

**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.



**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 4-46.

## Configuring iSLB

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

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

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

**Note**

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

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

This section covers the following topics:

- [About iSLB Configuration Limits, page 4-37](#)
- [iSLB Configuration Prerequisites, page 4-38](#)
- [About iSLB Initiators, page 4-38](#)
- [Configuring iSLB Using Device Manager, page 4-38](#)
- [Configuring iSLB Initiators, page 4-40](#)
- [About Load Balancing Using VRRP, page 4-45](#)
- [Configuring Load Balancing Using VRRP, page 4-47](#)
- [About iSLB Configuration Distribution Using CFS, page 4-47](#)
- [Distributing the iSLB Configuration Using CFS, page 4-48](#)

**Note**

Before configuring iSLB, you must enable iSCSI (see the [“Enabling iSCSI” section on page 4-5](#)).

**Note**

For iSLB, all switches in the fabric must be running Cisco MDS SAN-OS Release 2.1(1a) or later.

## About iSLB Configuration Limits

iSLB configuration has the following limits:

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

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

- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic is disrupted when any zoneset is activated.
- If IVR and iSLB features are enabled in the same fabric, you should have at least one switch in the fabric where both these features are enabled. Any zoning-related configuration and activation (for normal zones, IVR zones, or iSLB zones) must be performed on this switch. Otherwise, there may be traffic disruption in the fabric.

## iSLB Configuration Prerequisites

Perform the following prerequisite actions prior to configuring iSLB:

- Enable iSCSI (see the [“Enabling iSCSI” section on page 4-5](#) for more information).
- Configure the Gigabit Ethernet interfaces (see the [“Displaying Gigabit Ethernet Interface Statistics” section on page 6-14](#) and the [“Basic Gigabit Ethernet Configuration for IPv4” section on page 7-2](#)).
- Configure the VRRP groups (see the [“Configuring Load Balancing Using VRRP” section on page 4-47](#)).
- Configure and activate a zone set (see the *Cisco Fabric Manager Fabric Configuration Guide* for more information).
- Enable CFS distribution for iSLB (see the [“Enabling iSLB Configuration Distribution” section on page 4-48](#)).

## 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.
- 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.

## Configuring iSLB Using Device Manager


To configure iSLB using Device Manager, follow these steps:

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

**Step 1** Choose **IP > iSCSI iSLB**.

You see the iSCSI iSLB dialog box (see [Figure 4-32](#)).

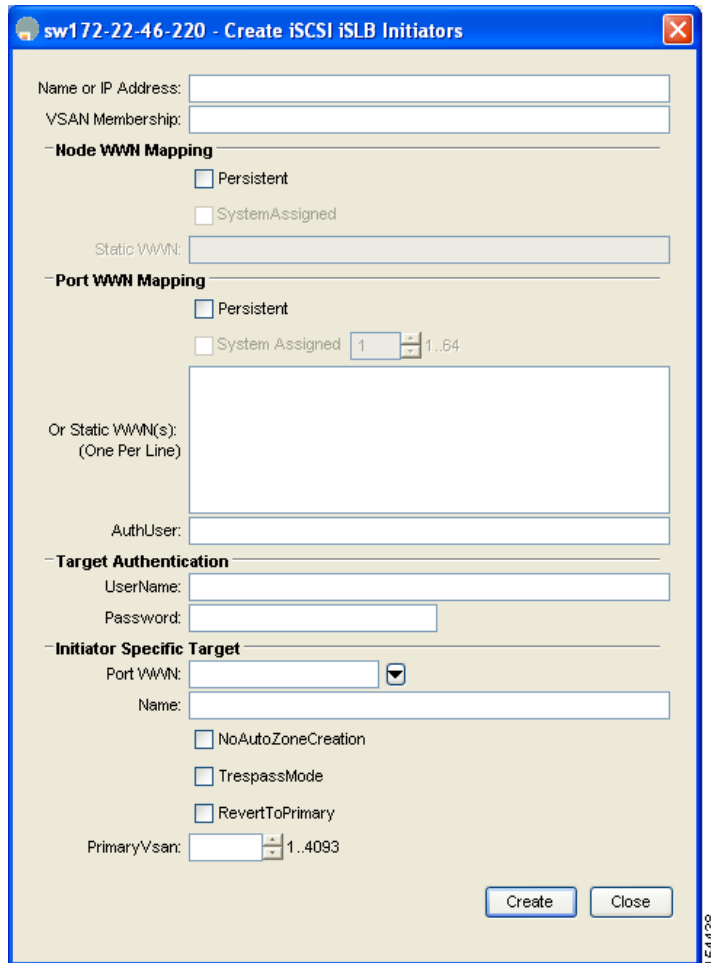
**Figure 4-32 iSCSI iSLB Dialog Box**



**Step 2** Click **Create** to create a new iSCSI iSLB initiator.

You see the Create iSCSI iSLB Initiators dialog box (see [Figure 4-33](#)).

**Figure 4-33 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.

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

Also see the “[Assigning VSAN Membership for iSLB Initiators](#)” section on page 4-41.

- Step 5** Check the **Persistent** check box to convert a dynamic nWWN to static for the iSLB initiator.  
Also see the “[Making the Dynamic iSLB Initiator WWN Mapping Static](#)” section on page 4-41.
- 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 iSLB Initiator WWN Mapping Static](#)” section on page 4-41.
- 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 4-44.
- Step 12** Fill in the Username and Password fields to configure iSLB initiator target CHAP authentication.  
Also see the “[Configuring iSLB Session Authentication](#)” section on page 4-44.
- 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 4-54.
- 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 Initiators

This section includes the following topics:

- [Assigning WWNs to iSLB Initiators](#), page 4-41
- [Making the Dynamic iSLB Initiator WWN Mapping Static](#), page 4-41
- [Configuring iSLB Target Access Mapping](#), page 4-63
- [Assigning VSAN Membership for iSLB Initiators](#), page 4-41
- [Configuring Metric for Load Balancing](#), page 4-42
- [Verifying iSLB Initiator Configuration](#), page 4-64



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

- [About Load Balancing Using VRRP](#), page 4-45
- [To configure additional iSLB initiator targets using Device Manager, follow these steps:](#), page 4-43
- [Configuring iSLB Session Authentication](#), page 4-69
- [Verifying iSLB Authentication Configuration](#), page 4-69

## 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 4-18.



### Tip

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

See the [“Configuring iSLB Using Device Manager”](#) procedure on page 4-38.

## 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 4-9).

You can convert a dynamic iSLB initiator to a static iSLB initiator and make its WWNs persistent



### Note

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).



### Note

Making the dynamic mapping for iSLB initiators static is the same as for iSCSI. See the [“Making the Dynamic iSLB Initiator WWN Mapping Static”](#) section on page 4-41.



### Note

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 4-38.

## 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.

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

For more information, see the *Cisco MDS 9000 Family NX-OS Fabric Manager Fabric Configuration Guide*.



**Note**

Specifying the iSLB initiator VSAN is the same as for an iSCSI initiator. See the [VSAN Membership for iSCSI, page 4-23](#).



**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 4-38](#).

## 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 4-45](#).

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 4-38](#).

## 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.

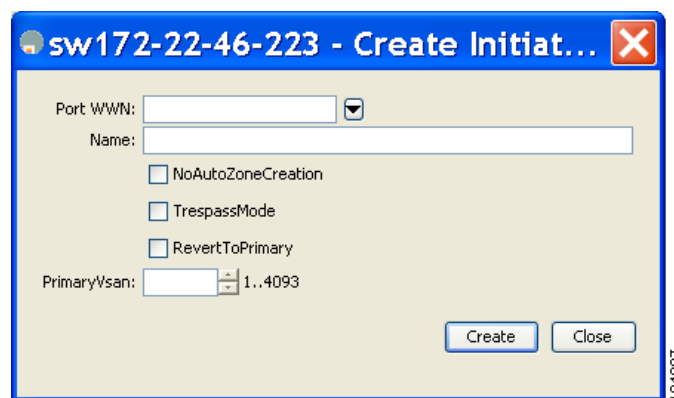
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

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

To configure additional iSLB initiator targets using Device Manager, follow these steps:

- 
- Step 1** Choose **IP > iSCSI iSLB**.  
You see the iSCSI iSLB dialog box (see [Figure 4-32](#)).
- 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 (see [Figure 4-34](#)).

**Figure 4-34 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 (see [Figure 4-33](#)).
- Step 7** (Optional) Check the **TrespassMode** check box. See the “LUN Trespass for Storage Port Failover” section on page 4-54.
- 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.
- 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 considerations:

- Auto-zoning of the initiator with the initiator targets is enabled by default.
- A zone set must be active in a VSAN for auto-zones to be created in that VSAN.

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

- 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.

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 4-38.

## Configuring iSLB Session Authentication

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

For CHAP user name or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see the *Cisco Fabric Manager 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 4-29.

### 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.

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 4-38.

### 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.

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.

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


See the “Configuring iSLB Using Device Manager” procedure on page 4-38.

## 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 4-35](#) shows an example of load balancing using iSLB.

**Figure 4-35 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.

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

**Note**

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

**Note**

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

**Tip**

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

**Caution**

A Gigabit Ethernet interface configured for iSLB can only be in one VRRP group because redirected sessions do not carry information about the VRRP IP address or group. This restriction allows 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:

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

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


## Configuring Load Balancing Using VRRP

You must first configure VRRP on the Gigabit Ethernet interfaces on the switch that connect to the IP network before configuring VRRP for iSLB. For information on how to configure VRRP on a Gigabit Ethernet interface, see the “[Virtual Router Redundancy Protocol](#)” section on page 5-10.

To configure VRRP load balancing using Device Manager, follow these steps:

- 
- Step 1** Choose **IP > iSCSI iSLB**.  
You see the iSCSI iSLB dialog box (see [Figure 4-32](#)).
  - 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 (see [Figure 4-36](#)).

**Figure 4-36** Create iSCSI iSLB VRRP Dialog Box



- Step 4** Set the Vrld 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.
- 

## 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 *Cisco Fabric Manager System Management Configuration Guide* for more information).

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

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

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

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

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



**Note**

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



**Note**

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

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.

## Distributing the iSLB Configuration Using CFS

This section contains the following:

- [Enabling iSLB Configuration Distribution, page 4-48](#)
- [Locking the Fabric, page 4-49](#)
- [Committing Changes to the Fabric, page 4-49](#)
- [Discarding Pending Changes, page 4-50](#)
- [Clearing a Fabric Lock, page 4-50](#)
- [CFS Merge Process, page 4-50](#)
- [iSLB CFS Merge Status Conflicts, page 4-51](#)

## Enabling iSLB Configuration Distribution

To enable CFS distribution of the iSLB configuration using Device Manager, follow these steps:

**Step 1**

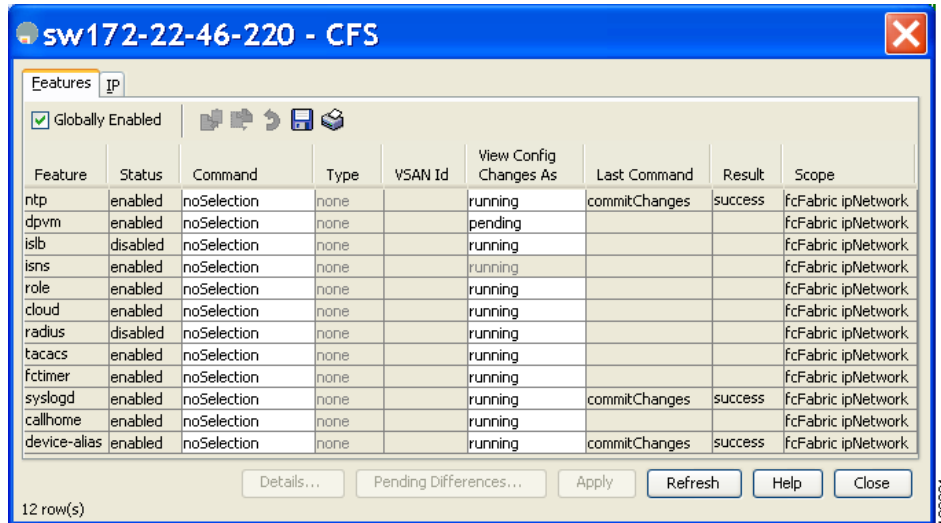
Choose **Admin > CFS**.

You see the CFS dialog box (see [Figure 4-37](#)).



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

**Figure 4-37 Enabling CFS in Device Manager**



**Step 2** Set the Command field to **enable** for the iSLB feature.

**Step 3** Click **Apply** to save this change.

## Locking the Fabric

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

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



### Note

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

## Committing Changes to the Fabric

To apply the pending iSLB configuration changes to the active configuration and to other MDS switches in the fabric, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the MDS switches throughout the fabric, the automatic zones are activated, and the fabric lock is released.

To commit iSLB configuration changes to other MDS switches in the fabric, activate iSLB automatic zones, and release the fabric lock using Device Manager, follow these steps:

**Step 1** Choose **Admin > CFS**.

You see the CFS Configuration dialog box (see [Figure 4-37](#)).

**Step 2** Set the Command field to **commit** for the iSLB feature.

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

**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.

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 (see [Figure 4-37](#)).
- Step 2** Set the Command field to **abort** for the iSLB feature.
- 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.



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

---

To release a fabric lock using Device Manager, follow these steps:

- 
- Step 1** Choose **Admin > CFS**.  
You see the CFS Configuration dialog box (see [Figure 4-37](#)).
- Step 2** Set the Command field to **clear** for the iSLB feature.
- Step 3** Click **Apply** to save this change.
- 

## 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.

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

- The iSLB initiator node names have no conflicts with iSCSI initiators on all the switches.

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

## 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 4-51](#)
- [Multiple IPS Ports Connected to the Same IP Network, page 4-54](#)
- [VRRP-Based High Availability, page 4-55](#)
- [Ethernet PortChannel-Based High Availability, page 4-56](#)

## Transparent Target Failover

The following high availability features are available for iSCSI configurations:

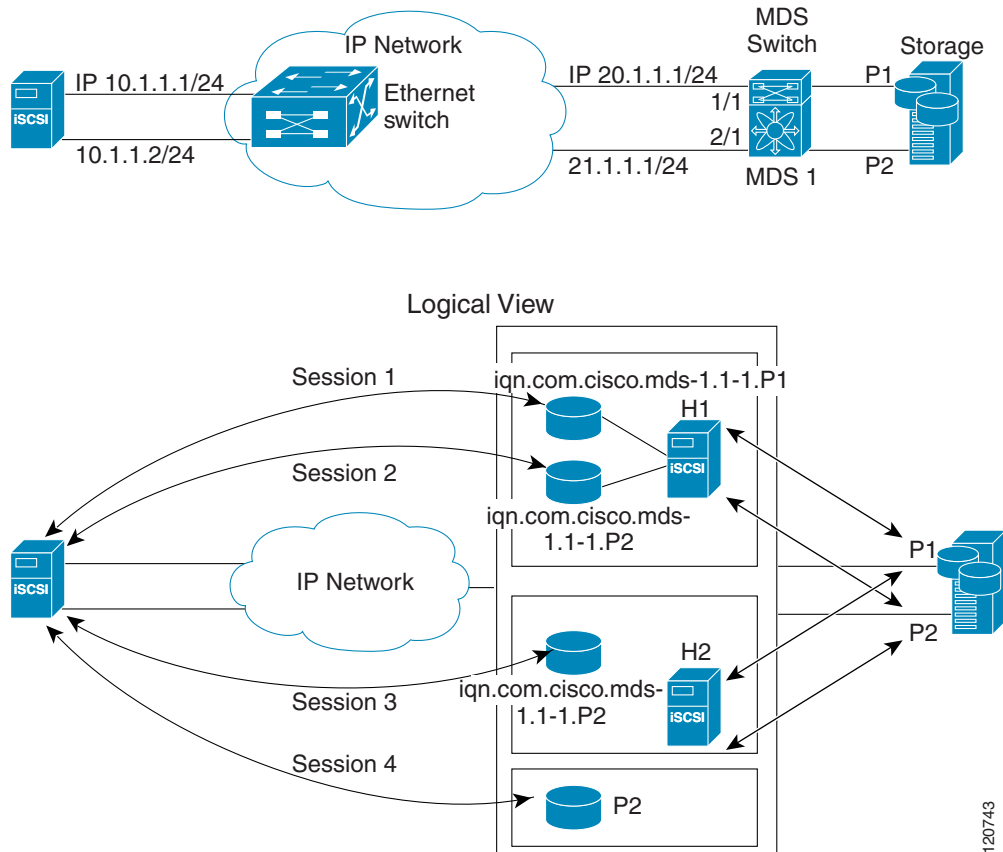
- iSCSI high availability with host running multi-path software—In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load balancing or failover across the different paths to access the storage.
- iSCSI high availability with host not having 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 4-38](#) 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 4-38 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 4-38 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.

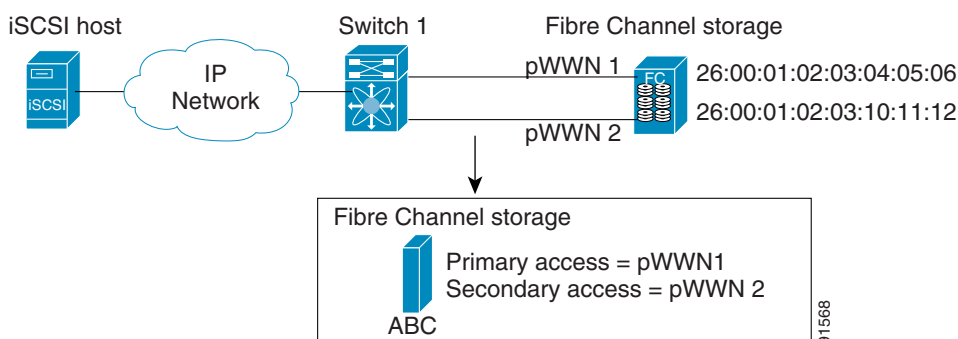
**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

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

- IPS ports support the VRRP feature (see [“Configuring VRRP for Gigabit Ethernet Interfaces” section on page 6-10](#)) to provide failover for IPS ports.
- IPS has transparent Fibre Channel target failover for iSCSI static virtual targets.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see [Figure 4-39](#)).

**Figure 4-39 Static Target Importing Through Two Fibre Channel Ports**



In [Figure 4-39](#), you can create an iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

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



**Tip**

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

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

To create a static iSCSI virtual target for the entire Fibre Channel target port using Device Manager, follow these steps:

- Step 1** Click **IP > iSCSI**.  
You see the iSCSI configuration (see [Figure 4-12](#)).
- Step 2** Click the **Targets** tab to display a list of existing iSCSI targets shown (see [Figure 4-13](#)).
- Step 3** Click **Create** to create an iSCSI target.  
You see the Create iSCSI Targets dialog box (see [Figure 4-15](#)).
- 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.

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


- 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 4-25.
- 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.
- Step 8** Click **Apply** to save this change.

## LUN Trespass for Storage Port Failover

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

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the trespass feature be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see [Figure 4-40](#)).

**Figure 4-40 Virtual Target with an Active Primary Port**



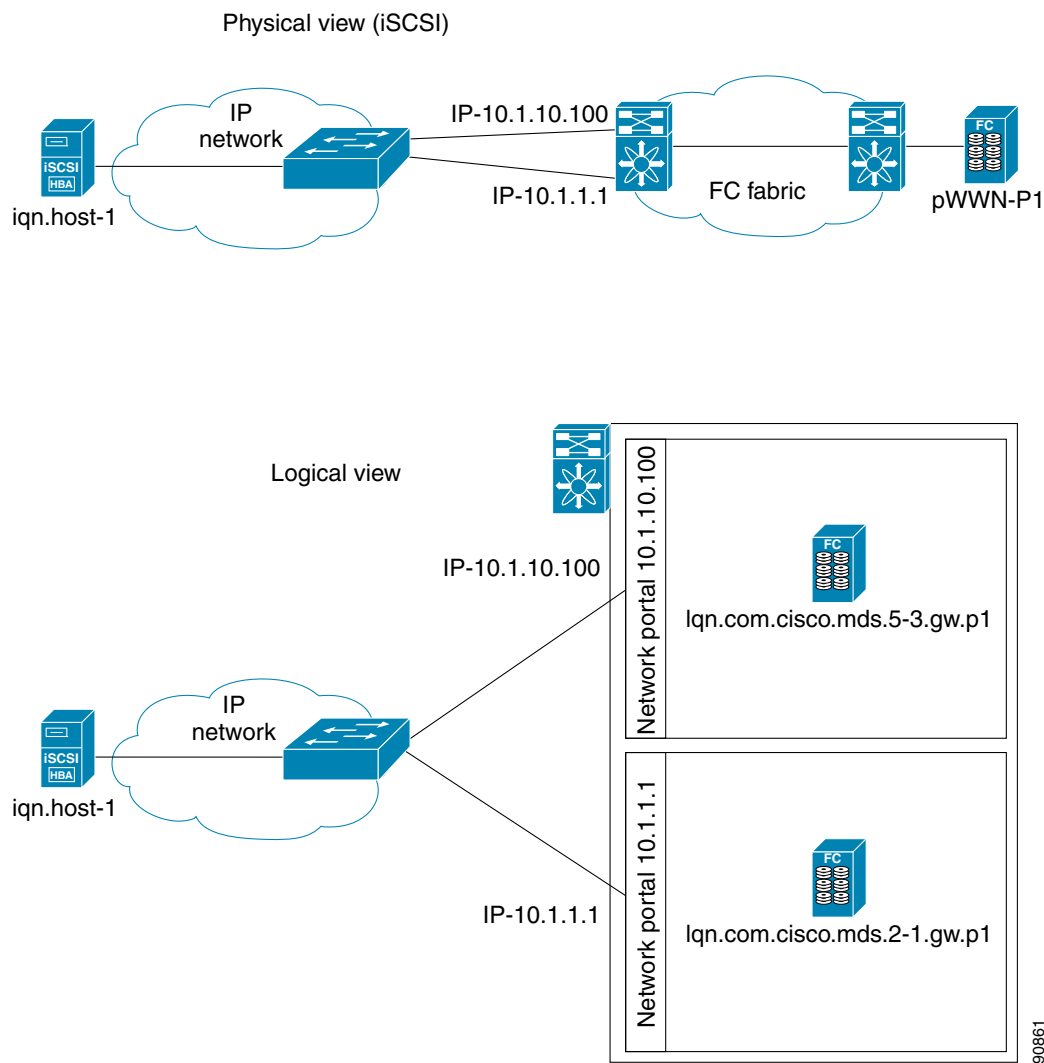
In Device Manager, choose **IP > iSCSI**, select the **Targets** tab, and check the **Trespass Mode** check box to enable the trespass feature for a static iSCSI virtual target.

## Multiple IPS Ports Connected to the Same IP Network

[Figure 4-41](#) provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

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

**Figure 4-41 Multiple Gigabit Ethernet Interfaces in the Same IP Network**



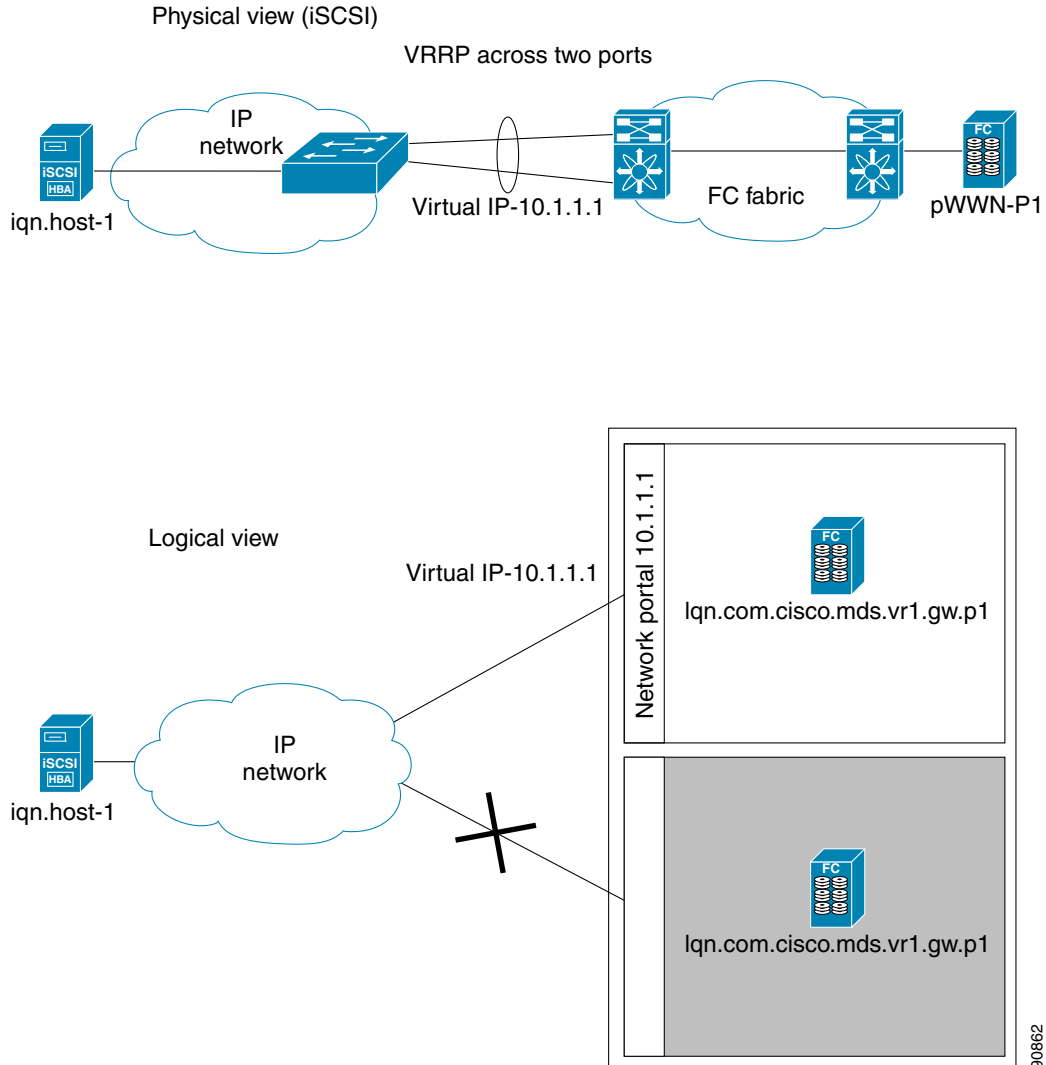
In [Figure 4-41](#), each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

## VRRP-Based High Availability

[Figure 4-42](#) provides an example of a VRRP-based high availability iSCSI configuration.

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

**Figure 4-42 VRRP-Based iSCSI High Availability**



In [Figure 4-42](#), 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 4-43](#) provides a sample Ethernet PortChannel-based high availability iSCSI configuration.



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

**Figure 4-43 Ethernet PortChannel-Based iSCSI High Availability**



In [Figure 4-43](#), each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.



**Note**

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

## iSCSI Authentication Setup Guidelines and Scenarios

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- [Configuring No Authentication, page 4-58](#)
- [Configuring CHAP with Local Password Database, page 4-58](#)
- [Configuring CHAP with External RADIUS Server, page 4-58](#)
- [iSCSI Transparent Mode Initiator, page 4-59](#)
- [Target Storage Device Requiring LUN Mapping, page 4-63](#)



**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 4-46](#).

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

## Configuring No Authentication

Set the iSCSI authentication method to **none** to configure a network with no authentication.

In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane. Then select the **Globals** tab and set the AuthMethod drop-down menu to **none** and click **Apply Changes**.

## Configuring CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

- 
- Step 1** Set the AAA authentication to use the local password database for the iSCSI protocol:
- In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.
  - Click the **Applications** tab in the Information pane.
  - 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:
- In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane.
  - Click the **Globals** tab in the Information pane.
  - Set the AuthMethod drop-down menu to **chap** and click **Apply Changes**.
- Step 3** Configure the user names and passwords for iSCSI users:
- In Device Manager, choose **Security > iSCSI**.
  - Set the Username, Password and Confirm Password fields.
  - Click **Create** to save these changes.
- Step 4** Verify the global iSCSI authentication setup:
- In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane.
  - Click the **Globals** tab in the Information pane.
- 

## Configuring CHAP with External RADIUS Server

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

- 
- Step 1** Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server:
- In Fabric Manager, 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 Fabric Manager, 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.

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

- d. Set the IP Type radio button to **ipv4** or **ipv6**.
  - e. 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:
- a. In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.
  - b. Select the **Server Groups** tab in the Information pane and click **Create Row**.
  - c. Set the Index field to a unique number.
  - d. Set the Protocol radio button to **radius**.
  - e. Set the Name field to the server group name.
  - f. 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.
- a. In Fabric Manager, choose **Switches > Security > AAA** in the Physical Attributes pane.
  - b. Click the **Applications** tab in the Information pane.
  - 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 Fabric Manager, 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 Fabric Manager, 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 Fabric Manager, 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.
- 

## iSCSI Transparent Mode Initiator

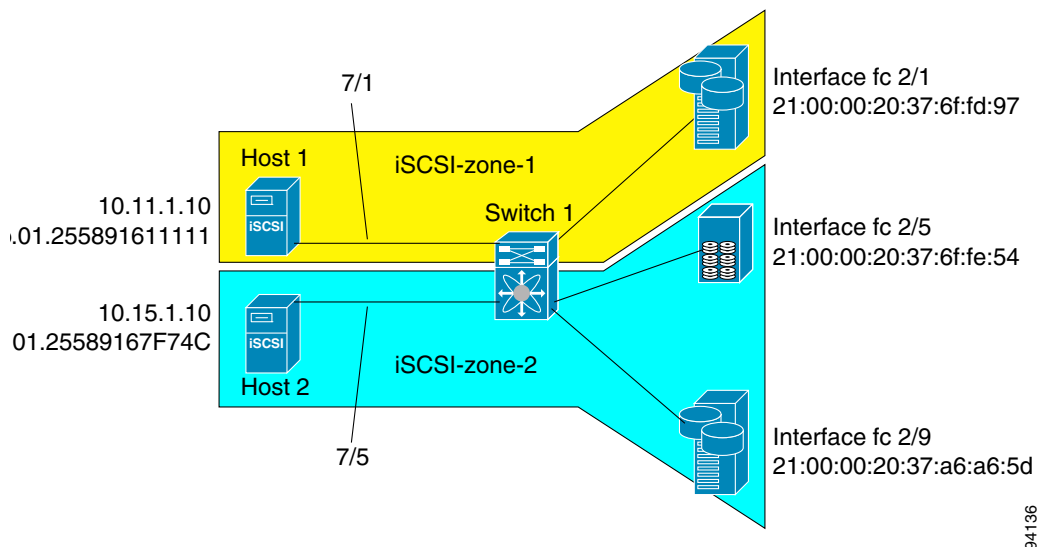
This scenario assumes the following configuration (see [Figure 4-44](#)):

- No LUN mapping or LUN masking or any other access control for hosts on the target device

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

- No iSCSI login authentication (that is, login authentication set to none)
- The topology is as follows:
  - iSCSI interface 7/1 is configured to identify initiators by IP address.
  - iSCSI interface 7/5 is configured to identify initiators by node name.
  - The iSCSI initiator host 1 with IPv4 address 10.11.1.10 and name iqn.1987-05.com.cisco:01.255891611111 connects to IPS port 7/1 is identified using IPv4 address (host 1 = 10.11.1.10).
  - The iSCSI initiator host 2 with IPv4 address 10.15.1.10 and node name iqn.1987-05.com.cisco:01.25589167f74c connects to IPS port 7/5.

**Figure 4-44 iSCSI Scenario 1**



94136

To configure scenario 1 (see [Figure 4-44](#)), follow these steps:

- 
- Step 1** Configure null authentication for all iSCSI hosts in Cisco MDS switches.
- a. In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane.
  - b. Select **none** from the AuthMethod drop-down menu in the Information pane.
  - c. 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.
- a. In Device Manager, click **IP > iSCSI**.
  - b. Click the **Targets** tab.
  - c. Check the **Dynamically Import FC Targets** check box.
  - d. Click **Apply**.
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IPv4 address and enable the interface.
- a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

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

- b. Select the **IP Address** tab in the Information pane and click **Create Row**.
- c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 1.
- d. Click **Create**.
- e. Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 1.
- f. 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.
- a. In Fabric Manager, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
  - b. Click the **iSCSI** tab in the Information pane.
  - c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
  - d. In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
  - e. Click the **iSCSI** tab.
  - f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.
  - g. Click **Apply**.
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with an IPv4 address and enable the interface.
- a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
  - b. Click the **IP Address** tab in the Information pane and click **Create Row**.
  - c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
  - d. Click **Create**.
  - e. Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
  - f. 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.
- a. In Fabric Manager, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
  - b. Click the **iSCSI** tab in the Information pane.
  - c. Select **name** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
  - d. In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
  - e. Click the **iSCSI** tab.
  - f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
  - g. Click **Apply**.



**Note** Host 1 is connected to this port.

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

**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 Fabric Manager, 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 Fabric Manager, 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)*



**Note** iSCSI interface is configured to identify all hosts based on node name.

- a. In Fabric Manager, 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 **zoneset-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 Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators
- Step 14** In Fabric Manager, 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

Sample scenario 2 assumes the following configuration (see [Figure 4-45](#)):

- 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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 4-45 iSCSI Scenario 2**



94137

To configure scenario 2 (see [Figure 4-45](#)), follow these steps:

- 
- Step 1** Configure null authentication for all iSCSI hosts.
- In Fabric Manager, 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 Fabric Manager, 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 Fabric Manager, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
  - Select the **iSCSI** tab in the Information pane.



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

- c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d. In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e. Click the **iSCSI** tab.
- f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 1.
- g. Click **Apply**.

**Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IPv4 address and enable the interface.

- a. In Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.
- b. Click the **IP Address** tab in the Information pane and click **Create Row**.
- c. Set the IP address and subnet mask for the Gigabit Ethernet interface in slot 7 port 5.
- d. Click **Create**.
- e. Select the **General** tab and select **up** from the Admin drop-down menu for the Gigabit Ethernet interface in slot 7 port 5.
- f. 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.

- a. In Fabric Manager, choose **Switches > Interfaces > FC Logical** in the Physical Attributes pane.
- b. Click the **iSCSI** tab in the Information pane.
- c. Select **ipaddress** from the Initiator ID Mode drop-down menu and click the **Apply Changes** icon.
- d. In Device Manager, choose **Interfaces > Ethernet and iSCSI**.
- e. Click the **iSCSI** tab.
- f. Select **up** from the Admin drop-down menu for the iSCSI interface in slot 7 port 5.
- g. Click **Apply**.

**Step 7** Configure for static pWWN and nWWN for host 1.

- a. In Device Manager, choose **IP > iSCSI**.
- b. Click the **Initiators** tab.
- c. Check the **Node Address Persistent** and **Node Address System-assigned** check boxes the Host 1 iSCSI initiator.
- d. Click **Apply**.

**Step 8** Configure for static pWWN for Host 2.

- a. In Device Manager, Choose **IP > iSCSI**.
- b. Click the **Initiators** tab.
- c. Right-click on the Host 2 iSCSI initiator and click Edit pWWN.
- d. 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.

- a. In Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane.

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

- b. Click the **Initiators** tab.

**Step 10** Create a zone for Host 1 and the iSCSI target in VSAN 1.



**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 Fabric Manager, 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 Fabric Manager, 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 Fabric Manager, 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**.

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

- 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 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 Fabric Manager, 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 **zoneset-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 Fabric Manager, choose **End Devices > iSCSI** in the Physical Attributes pane to verify the details of the two iSCSI initiators.

**Step 17** In Fabric Manager, 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.

---

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

## iSNS

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

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

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

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

This section includes the following topics:

- [About iSNS Client Functionality, page 4-68](#)
- [Creating an iSNS Client Profile, page 4-69](#)
- [About iSNS Server Functionality, page 4-70](#)
- [Configuring iSNS Servers, page 4-72](#)

## 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 4-8](#) for more details on how iSCSI imports Fibre Channel targets.

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

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

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

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


**Note**

The iSNS client is not supported on a VRRP interface.

## Creating an iSNS Client Profile

To create an iSNS profile using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** in the Physical Attributes pane.
- You see the iSCSI configuration in the Information pane (see [Figure 4-12](#)).
- Step 2** Select the **iSNS** tab.
- Step 3** You see the iSNS profiles configured (see [Figure 4-46](#)).

**Figure 4-46 iSNS Profiles in Fabric Manager**

Switch	Feature	Status	Command	LastCommand	Result
sw172-22-46-220	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-223	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-233	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-224	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-182	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-223	iSNS-server	enabled	noSelection	noSelection	none
sw172-22-46-221	iscsi	disabled	noSelection	noSelection	none
sw172-22-46-222	iscsi	enabled	noSelection	noSelection	none
sw172-22-46-233	iSNS-server	enabled	noSelection	noSelection	none

- 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 delete an iSNS profile using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** from the Physical Attributes pane.
- You see the iSCSI configuration in the Information pane (see [Figure 4-12](#)).
- Step 2** Select the **iSNS** tab.
- You see the iSNS profiles configured (see [Figure 4-46](#)).

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

**Step 3** Right-click on the profile that you want to delete and click the **Delete Row** icon.

To tag a profile to an interface using Fabric Manager, follow these steps:

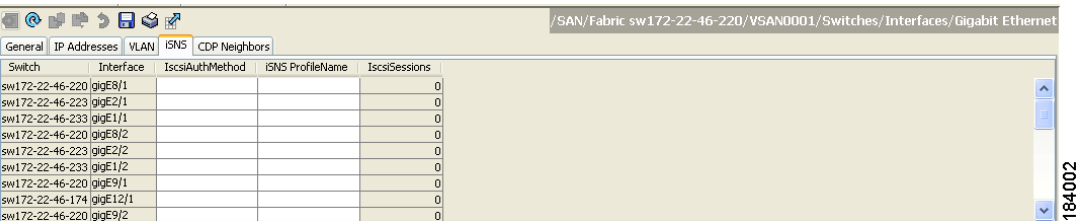
**Step 1** Choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane.

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 (see [Figure 4-47](#)).

**Figure 4-47 iSNS Profiles in Fabric Manager**



Switch	Interface	IscliAuthMethod	iSNS ProfileName	IscliSessions
sw172-22-46-220	giqE8/1			0
sw172-22-46-223	giqE2/1			0
sw172-22-46-233	giqE1/1			0
sw172-22-46-220	giqE8/2			0
sw172-22-46-223	giqE2/2			0
sw172-22-46-233	giqE1/2			0
sw172-22-46-220	giqE9/1			0
sw172-22-46-174	giqE12/1			0
sw172-22-46-220	giqE9/2			0

**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 using Fabric Manager, 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 (see [Figure 4-47](#)).

**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.

# 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.


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

- Provides a notification mechanism for registered iSNS clients to receive change notifications on the status change of other iSNS clients.
- Provides a single access control configuration for both Fibre Channel and iSCSI devices.
- Discovers iSCSI targets that do not have direct IP connectivity to the iSCSI initiators.

## Example Scenario

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. [Figure 4-48](#) provides an example of this scenario.

**Figure 4-48 Using iSNS Servers in the Cisco MDS Environment**



In [Figure 4-48](#), 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).

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

7. If the initiator chooses to log in to SW-1 and later that port becomes inaccessible (for example, Gigabitethernet 2/1 goes down), the initiator has the choice to move to connect to target P1 through port Gigabitethernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

## Configuring iSNS Servers

This section describe how to configure an iSNS server on a Cisco MDS 9000 Family switch.

This section includes the following topics:

- [Enabling the iSNS Server, page 4-72](#)
- [iSNS Configuration Distribution, page 4-72](#)
- [Configuring the ESI Retry Count, page 4-119](#)
- [Configuring the Registration Period, page 4-73](#)
- [iSNS Client Registration and Deregistration, page 4-73](#)
- [Target Discovery, page 4-74](#)
- [Verifying the iSNS Server Configuration, page 4-120](#)

### Enabling the iSNS Server

Before the iSNS server feature can be enabled, iSCSI must be enabled (see the “[Enabling iSCSI](#)” section on [page 4-5](#)). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

To enable the iSNS server using Fabric Manager, 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 5-13](#)).

### 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 *Cisco Fabric Manager System Management Configuration Guide*.



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

To enable iSNS configuration distribution using Fabric Manager, 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.

## Configuring the Registration Period

The iSNS client specifies the registration period with the iSNS Server. The iSNS Server keeps the registration active until the end of this period. If there are no commands from the iSNS client during this period, then the iSNS server removes the client registration from its database.

If the iSNS client does not specify a registration period, the iSNS server assumes a default value of 0, which keeps the registration active indefinitely. You can also manually configure the registration period on the MDS iSNS Server.

To configure the registration period on an iSNS Server using Fabric Manager, 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.
- 

## iSNS Client Registration and Deregistration

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).

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

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.

## iSNS Cloud Discovery

You can configure iSNS cloud discovery to automate the process of discovering iSNS servers in the IP network.

This section includes the following topics:

- [About Cloud Discovery, page 4-74](#)
- [Configuring iSNS Cloud Discovery, page 4-75](#)

## 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 [fm-docfeedback@cisco.com](mailto:fm-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.

## Configuring iSNS Cloud Discovery

This section describes how to configure iSNS cloud discovery and includes the following topics:

- [Enabling iSNS Cloud Discovery, page 4-75](#)
- [Initiating On-Demand iSNS Cloud Discovery, page 4-129](#)
- [Configuring Automatic iSNS Cloud Discovery, page 4-76](#)
- [Configuring iSNS Cloud Discovery Distribution, page 4-129](#)

## Enabling iSNS Cloud Discovery

To enable iSNS cloud discovery using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>End Devices &gt; iSNS</b> .<br><br>You see the iSNS configuration in the Information pane.                   |
| <b>Step 2</b> | Click the <b>Control</b> tab and select <b>enable</b> from the Command drop-down menu for the cloud discovery feature. |

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

**Step 3** Click the **Apply Changes** icon to save this change.

---

## Initiating On-Demand iSNS Cloud Discovery

To initiate on-demand iSNS cloud discovery using Fabric Manager, 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.
- 

## Configuring Automatic iSNS Cloud Discovery

To configure automatic iSNS cloud discovery using Fabric Manager, 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

To configure iSNS cloud discovery CFS distribution using Fabric Manager, 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.
- 

## Default Settings

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

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

**Table 4-2 Default iSCSI Parameters**

Parameters	Default
Number of TCP connections	One per iSCSI session
<b>minimum-retransmit-time</b>	300 msec
<b>keepalive-timeout</b>	60 seconds
<b>max-retransmissions</b>	4 retransmissions
PMTU discovery	Enabled
<b>pmtu-enable reset-timeout</b>	3600 sec
SACK	Enabled
<b>max-bandwidth</b>	1 Gbps
<b>min-available-bandwidth</b>	70 Mbps
<b>round-trip-time</b>	1 msec
Buffer size	4096 KB
Control TCP and data connection	No packets are transmitted
TCP congestion window monitoring	Enabled
Burst size	50 KB
Jitter	500 microseconds
TCP connection mode	Active mode is enabled
Fibre Channel targets to iSCSI	Not imported
Advertising iSCSI target	Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces
iSCSI hosts mapping to virtual Fibre Channel hosts	Dynamic mapping
Dynamic iSCSI initiators	Members of the VSAN 1
Identifying initiators	iSCSI node names
Advertising static virtual targets	No initiators are allowed to access a virtual target (unless explicitly configured)
iSCSI login authentication	CHAP or none authentication mechanism
<b>revert-primary-port</b>	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.

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

**Table 4-3**      **Default iSLB Parameters**

Parameters	Default
Fabric distribution	Disabled
Load balancing metric	1000



## CHAPTER 6

# Configuring IP Storage

---

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



### Note

FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.

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

This chapter includes the following sections:

- [IP Storage Services Modules, page 6-1](#)
- [Supported Hardware, page 6-4](#)
- [Configuring Gigabit Ethernet Interfaces for IPv4, page 6-4](#)
- [IPS Module Core Dumps, page 6-5](#)
- [Configuring Gigabit Ethernet High Availability, page 6-9](#)
- [Default Settings, page 6-11](#)

## IP Storage Services Modules

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.
- **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.

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


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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 6-2 iSCSI Scenarios**



## Module Status Verification

To verify the status of the module using Fabric Manager, 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.
- 

## 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.



### Note

The IPS-8 (DS-X9308-SMIP) and IPS-4 (DS-X9304-SMIP) do not support NX-OS 4.x or above." on both the 4.x and 5.x config guides.

***Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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

The IPS-8 (DS-X9308-SMIP) and IPS-4 (DS-X9304-SMIP) do not support NX-OS 4.x or above." on both the 4.x and 5.x config guides.



### 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*).

## Configuring Gigabit Ethernet Interfaces for IPv4

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured.

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.”](#)

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

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.



**Note**

The Gigabit Ethernet interfaces on the MPS-14/2 module do not support EtherChannel.



**Note**

To configure IPv6 on a Gigabit Ethernet interface, see the *Cisco Fabric Manager Security Configuration Guide*.




**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.

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

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module because of its large space requirement. They are copied directly to an external TFTP server.

Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps). To configure the Gigabit Ethernet interface for the scenario in [Figure 6-3](#), follow these steps:

- 
- Step 1** From Fabric Manager, 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 Fabric Manager, 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 Fabric Manager, 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 Interface Descriptions

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the switch port description for any interface.

## Configuring Beacon Mode

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the beacon mode for any interface.

## Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

## 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.

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



**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.

## 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.

## 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.

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.

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

**Table 6-1 Subnet Requirements for Interfaces (continued)**

Interface 1	Interface 2	Same Subnet Allowed	Notes
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.

## Gigabit Ethernet IPv4-ACL Guidelines



**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:

- 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.

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

- 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.

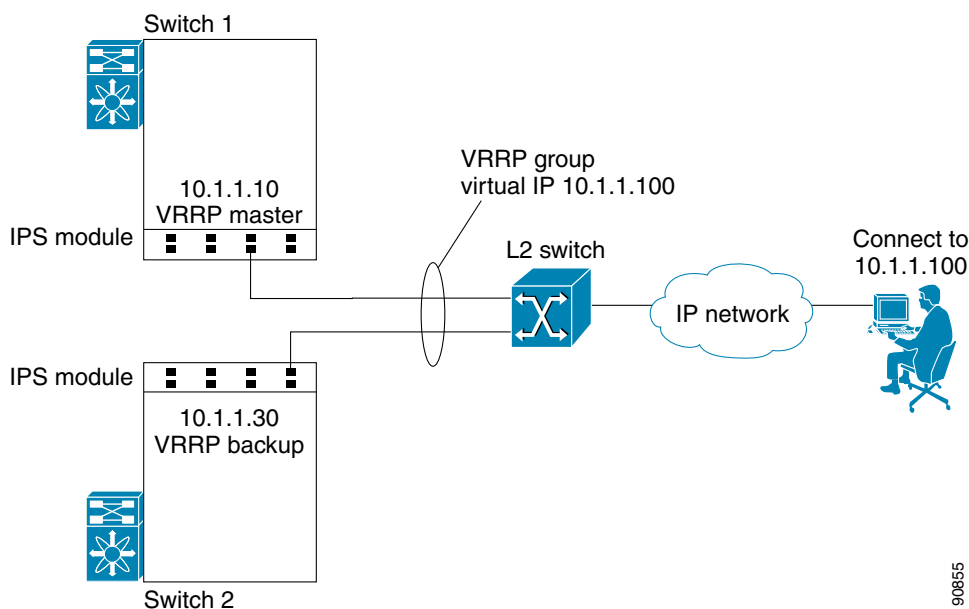
## Configuring Gigabit Ethernet High Availability

Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

### VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address failover protection to an alternate Gigabit Ethernet interface so the IP address is always available (see [Figure 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

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



**Note**

You can configure no more than seven VRRP groups, both IPv4 and IPv6, on a Gigabit Ethernet interface, including the main interface and all subinterfaces.

## Configuring VRRP for Gigabit Ethernet Interfaces



**Note**

If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.



**Note**

The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual 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.

## 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.



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

**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.

## Configuring Ethernet PortChannels

The PortChannel configuration specified in the *Cisco Fabric Manager Interfaces Configuration Guide* also applies to Ethernet PortChannel configurations.

## Configuring CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of Layer 3 protocols. 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*.

## 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

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



## CHAPTER 7

# Configuring IPv4 for Gigabit Ethernet Interfaces

Cisco MDS 9000 Family supports IP version 4 (IPv4) on Gigabit Ethernet interfaces. This chapter describes how to configure IPv4 addresses and other IPv4 features.

This chapter includes the following topics:

- [About IPv4, page 7-1](#)
- [Basic Gigabit Ethernet Configuration for IPv4, page 7-2](#)
- [VLANs, page 7-4](#)
- [IPv4-ACLs, page 7-6](#)
- [Default Settings, page 7-7](#)

## 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.”](#)

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.

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

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 8-11](#).




**Tip**

Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port. They should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

## Basic Gigabit Ethernet Configuration for IPv4

Figure 7-1 shows an example of a basic Gigabit Ethernet IP version 4 (IPv4) configuration.

**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.

## Configuring Gigabit Ethernet Interface

To configure the Gigabit Ethernet interface using Fabric Manager, 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**.

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

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.
- 

This section includes the following topics:

- [Configuring Interface Descriptions, page 7-3](#)
- [Configuring Beacon Mode, page 7-3](#)
- [Configuring Autonegotiation, page 7-3](#)
- [Configuring the MTU Frame Size, page 7-4](#)
- [Configuring Promiscuous Mode, page 7-4](#)

## Configuring Interface Descriptions

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the switch port description for any interface.

## Configuring Beacon Mode

See the *Cisco Fabric Manager Interfaces Configuration Guide* for details on configuring the beacon mode for any interface.

## Configuring Autonegotiation

By default, autonegotiation is enabled all Gigabit Ethernet interface. You can enable or disable autonegotiation for a specified Gigabit Ethernet interface. When autonegotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the autonegotiation feature.

To configure autonegotiation using Fabric Manager, 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**.
-

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

## 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.

To configure the MTU frame size using Fabric Manager, 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**.
- 

## Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

To configure the promiscuous mode using Fabric Manager, 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**.
- 

## VLANs

This section describes virtual LAN (VLAN) support in Cisco MDS NX-OS and includes the following topics:

- [About VLANs for Gigabit Ethernet, page 7-5](#)
- [Configuring the VLAN Subinterface, page 7-5](#)
- [Interface Subnet Requirements, page 7-5](#)

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

## About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.



### Note

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name:

slot-number / port-number.VLAN-ID

## Configuring the VLAN Subinterface

To configure a VLAN subinterface (VLAN ID) using Device Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Select <b>Interface &gt; Ethernet and iSCSI</b> .                          |
| <b>Step 2</b> | Click the <b>Sub Interfaces</b> tab.                                       |
| <b>Step 3</b> | Select the Gigabit Ethernet subinterface on which 802.1Q should be used.   |
| <b>Step 4</b> | Click the <b>Edit IP Address</b> button.                                   |
| <b>Step 5</b> | Enter the IPv4 address and subnet mask for the Gigabit Ethernet interface. |
| <b>Step 6</b> | Click <b>Create</b> to save the changes or you may click <b>Close</b> .    |
- 

## 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](#)).

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

**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.

## IPv4-ACLs

This section describes the guidelines for IPv4 access control lists (IPv4-ACLs) and how to apply them to Gigabit Ethernet interfaces.



**Note**

For information on creating IPv4-ACLs, see the *Cisco Fabric Manager Security Configuration Guide*.

## Gigabit Ethernet IPv4-ACL Guidelines

Follow these guidelines when configuring IPv4-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP) or Internet Control Message Protocol (ICMP).



**Note**

Other protocols such as User Datagram Protocol (UDP) and HTTP are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IPv4-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
  - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
  - The **established** option is ignored when you apply IPv4-ACLs containing this option to Gigabit Ethernet interfaces.



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

- 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. For information on configuring IPv4-ACLs

*Cisco Fabric Manager Security Configuration Guide*

## 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

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



## CHAPTER 8

# 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.

This chapter includes the following sections:

- [About IPv6, page 8-1](#)
- [Configuring Basic Connectivity for IPv6, page 8-11](#)
- [Configuring IPv6 Static Routes, page 8-13](#)
- [Gigabit Ethernet IPv6-ACL Guidelines, page 8-14](#)
- [Transitioning from IPv4 to IPv6, page 8-15](#)
- [Default Settings, page 8-15](#)



### 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.”](#)

## 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).
- Provides simpler autoconfiguration of addresses.

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

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](#)

## 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 [fm-docfeedback@cisco.com](mailto:fm-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	FF01:0:0:0:0:0:0:101	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 [fm-docfeedback@cisco.com](mailto:fm-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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

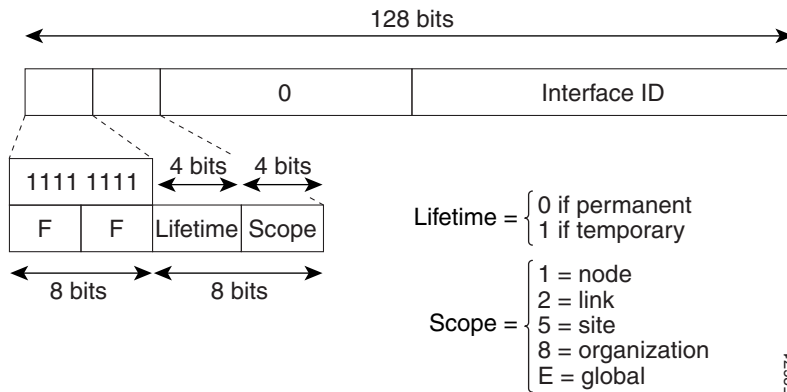
**Figure 8-3 Link-Local Address Format**



## IPv6 Address Type: Multicast

An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (1111 1111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. Figure 8-4 shows the format of the IPv6 multicast address.

**Figure 8-4 IPv6 Multicast Address Format**



IPv6 hosts are required to join (receive packets destined for) the following multicast groups:


- All-node multicast group FF02::1.
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 concatenated with the low-order 24 bit of the unicast address.

The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast address. IPv6 nodes must join the associated solicited-node multicast group for every unicast address to which it is assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6

**Send documentation comments to [fm-docfeedback@cisco.com](mailto:fm-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

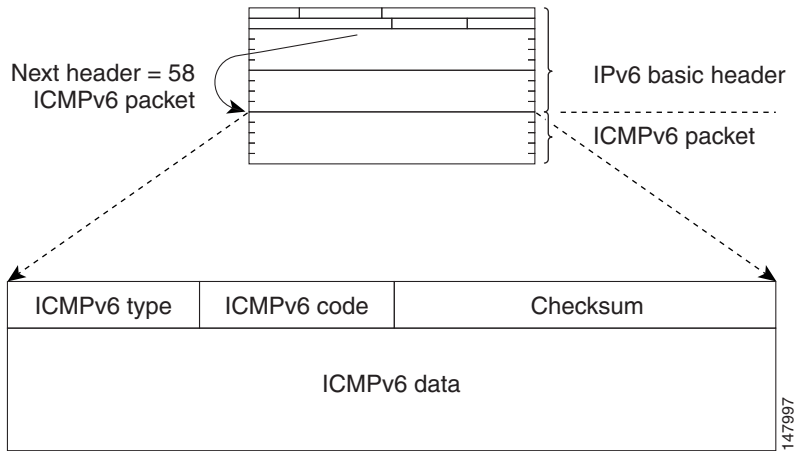
Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4—ICMP generates error messages such as ICMP destination unreachable messages, and informational messages such as ICMP echo request and reply messages. Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery 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 [fm-docfeedback@cisco.com](mailto:fm-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

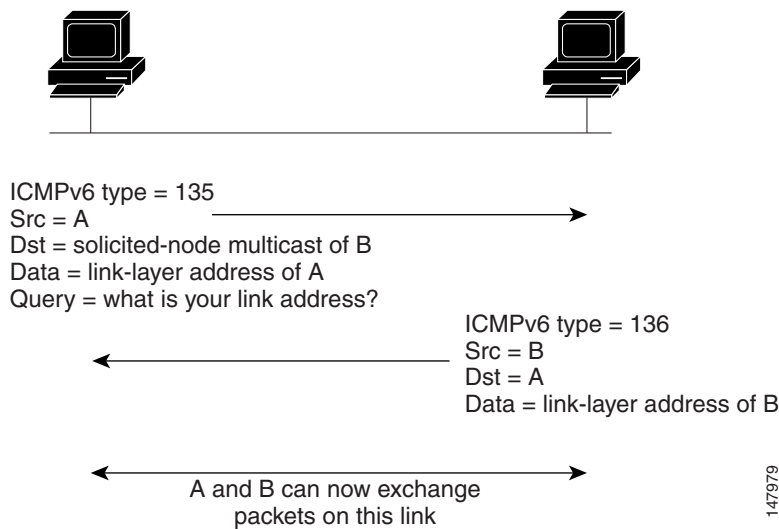
The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighboring routers.

### 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 [fm-docfeedback@cisco.com](mailto:fm-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 [fm-docfeedback@cisco.com](mailto:fm-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 [fm-docfeedback@cisco.com](mailto:fm-docfeedback@cisco.com)**

**Figure 8-8 IPv6 Stateless Autoconfiguration**

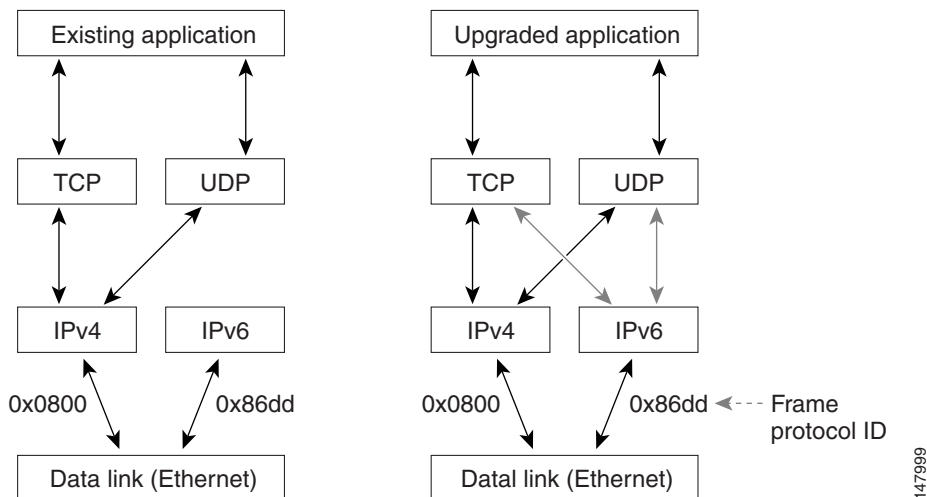


A node on the link can automatically configure site-local and global IPv6 addresses by appending its interface identifier (64 bits) to the prefixes (64 bits) included in the RA messages. The resulting 128-bit IPv6 addresses configured by the node are then subjected to duplicate address detection to ensure 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 [fm-docfeedback@cisco.com](mailto:fm-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**



## Configuring Basic Connectivity for IPv6

The tasks in this section explain how to implement IPv6 basic connectivity. Each task in the list is identified as either required or optional. This section includes the following topics:

- [Configuring IPv6 Addressing and Enabling IPv6 Routing, page 8-11](#)
- [Configuring IPv4 and IPv6 Protocol Addresses, page 8-13](#)

## Configuring 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: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`

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

This task explains how to assign IPv6 addresses to individual router interfaces and enable the processing of IPv6 traffic. By default, IPv6 addresses are not configured and IPv6 processing is disabled.

You can configure IPv6 addresses on the following interface types:

- Gigabit Ethernet
- Management
- VLAN (Gigabit Ethernet subinterface)
- VSAN



### Note

The IPv6 address 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 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 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.

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 (see [Figure 8-11](#)).

**Figure 8-11 Gigabit Ethernet Configuration in Device Manager**



**Step 2** Click the IP Address that you want to configure and click **Edit IP Address**.

You see the IP Address dialog box.


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

- Step 3** Click **Create** and set the IP Address/Mask field, using the IPv6 format (for example, 2001:0DB8:800:200C::417A/64).
- 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. (see [Figure 8-11](#)).

**Figure 8-12 IP Routing Configuration in Device Manager**



- 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.

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.
- 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.

## Configuring IPv6 Static Routes

Cisco MDS NX-OS supports static routes for IPv6. This section includes the following topics:

- [Configuring a IPv6 Static Route, page 8-14](#)

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

## Configuring a IPv6 Static Route

You must manually configure IPv6 static routes and define an explicit path between two networking devices. IPv6 static routes are not automatically updated and must be manually reconfigured if the network topology changes.

To configure a IPv6 static route using Device Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>IP &gt; Routing</b> .<br>You see the IP Routing Configuration dialog box.             |
| <b>Step 2</b> | Click <b>Create</b> .<br>You see the Create IP Route dialog box.                                |
| <b>Step 3</b> | Set the Dest field to the IPv6 destination address.   |
| <b>Step 4</b> | Set the Mask field to the IPv6 subnet mask.   |
| <b>Step 5</b> | Set the Gateway field to the IPv6 default gateway.  |
| <b>Step 6</b> | (Optional) Set the Metric field to the desired route metric.                                    |
| <b>Step 7</b> | Select the interface from the Interface drop-down menu.   |
| <b>Step 8</b> | Click <b>Create</b> to save these changes or click <b>Close</b> to discard any unsaved changes. |
- 

## Gigabit Ethernet IPv6-ACL Guidelines



### Tip

If IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See the *Cisco Fabric Manager 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.

- 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.



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

- 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 *Cisco Fabric Manager Security Configuration Guide* for information on applying IPv6-ACLs to an interface.

## 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*.

## 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

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



## INDEX

---

### A

AAA authentication  
    configuring [4-29, 4-30](#)  
access control  
    enforcing iSCSI  
        enforcing access control [4-28](#)  
    iSCSI [4-27](#)  
access control zoning based access control iSCSI  
    zoning based access control [4-28](#)  
ACL based access control  
    configuring for iSCSI [4-27](#)  
ACLs  
    configuring for iSCSI [4-27](#)  
advertised interfaces [4-13](#)  
advertisement packets  
    setting time intervals [5-13](#)  
authentication  
    CHAP option [4-58](#)  
    configuring local with Device Manager [4-31](#)  
    iSCSI setup [4-57](#)  
    local [4-31](#)  
    MD5 [5-14](#)  
    mechanism [4-30](#)  
    mutual CHAP authentication [4-31](#)  
    restricting iSLB initiator authentication  
        restricting iSLB  
            restricting iSLB initiators [4-44](#)  
    simple text [5-14](#)  
    See also MD5 authentication  
    See also simple text authentication  
autogenerated iSCSI target iSCSI  
    autogenerated target [4-29](#)

auto-negotiation  
    configuring Gigabit Ethernet interfaces [6-5, 7-3](#)

---

### B

B ports  
    configuring [2-24](#)  
    interoperability mode [2-22](#)  
    SAN extenders [2-23](#)  
bridge ports. See B ports  
buffer sizes  
    configuring in FCIP profiles [2-20](#)

---

### C

CDP  
    configuring  
CFS  
    iSLB config distribution [4-47](#)  
CHAP authentication [4-29, 4-44, 4-58](#)  
CHAP challenge [4-31](#)  
CHAP response [4-31](#)  
CHAP user name [4-31](#)  
Cisco Discovery Protocol. See CDP  
Cisco Transport Controller. See CTC  
cloud discovery. See iSNS cloud discovery  
congestion window monitoring. See CWM  
CTC  
    description [2-17](#)  
    launching [2-17](#)  
Cut-through routing mode [4-34](#)  
cut-thru routing mode [4-36](#)  
CWM

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

configuring in FCIP profiles [2-19](#)

## D

default gateways. See IPv4 default gateways

default networks. See IPv4 default networks

differentiated services code point. See DSCP

### DNS

default settings [5-15](#)

### DNS servers

configuring [5-14](#)

### drivers

iSCSI [4-2](#)

### DSCP

configuring [2-25](#)

### dynamic iSCSI initiator

converting [4-41](#)

convert to static iSCSI

convert dynamic initiator to static [4-20](#)

dynamic mapping [4-9, 4-41](#)

### dynamic mapping iSCSI

dynamic mapping iSCSI

static mapping static mapping [4-8](#)

## E

### ELP

verifying using Device Manager (procedure) [2-16](#)

entity status inquiry. See ESI

### E ports

configuring [2-25](#)

### ESI

non-resp threshold [4-73](#)

ESI retry count [4-73](#)

### Ethernet PortChannels

adding Gigabit Ethernet interfaces [6-10](#)

configuring [6-10](#)

description [6-9](#)

iSCSI [4-56](#)

redundancy [2-6](#)

explicit fabric logout [4-17](#)

Extended Link Protocol. See ELP

external RADIUS server

CHAP [4-58](#)

external RADIUS servers

CHAP [4-58](#)

## F

### fabric lock

releasing [4-50](#)

### FCIP [4-1](#)

advanced features [2-26](#)

checking trunk status (procedure) [2-17](#)

compression [2-33](#)

configuring [2-7](#)

configuring using FCIP Wizard [2-8 to 2-15](#)

default parameters [2-34](#)

discarding packets [2-22](#)

enabling [2-8](#)

Gigabit Ethernet ports [6-4, 7-1](#)

high availability [2-4](#)

IPS modules [2-2](#)

IP storage services support [6-1](#)

link failures [2-5](#)

MPS-14/2 module [2-2](#)

tape acceleration [2-28](#)

time stamps [2-22](#)

VE ports [2-2](#)

verifying ELP (procedure) [2-16](#)

verifying interfaces (procedure) [2-16](#)

virtual ISLs [2-2](#)

VRRP [2-6](#)

write acceleration [2-26](#)

### FCIP compression

configuring (procedure) [2-12](#)

description [2-33](#)

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

## FCIP interfaces

- configuring advanced features [2-20 to 2-25](#)
- configuring peers [2-20](#)
- configuring QoS [2-25](#)
- creating [2-20](#)
- parameters [2-4](#)

## FCIP links

- B port interoperability mode [2-22](#)
- configuring [2-15](#)
- configuring peers [2-20](#)
- configuring QoS [2-25](#)
- creating [2-16](#)
- description [2-3](#)
- endpoints [2-3](#)
- initiating IP connections [2-22](#)
- TCP connections [2-3](#)

## FCIP peers

- configuring IP addresses [2-21](#)

## FCIP profiles

- configuring TCP parameters [2-17](#)
- creating [2-15](#)
- description [2-4](#)

## FCIP tape acceleration

- configuring [2-33](#)
- description [2-28 to 2-33](#)

## FCIP TCP parameters

- configuring buffer size [2-20](#)
- configuring CWM [2-19](#)
- configuring keepalive timeouts [2-18](#)
- configuring maximum jitter [2-20](#)
- configuring maximum retransmissions [2-18](#)
- configuring minimum retransmit timeouts [2-18](#)
- configuring PMTUs [2-18](#)
- configuring SACKs [2-18](#)
- configuring window management [2-19](#)

## FCIP write acceleration

- configuring [2-28](#)
- configuring (procedure) [2-12](#)
- description [2-26](#)

## FC Logical Interface Tables [4-22](#)

## FCP

- routing requests [4-3](#)

## Fibre Channel [4-1](#)

- iSCSI targets [4-8 to 4-14](#)

## Fibre Channel over IP. See FCIP

## Fibre Channel targets

- dynamic importing [4-10](#)
- dynamic mapping [4-10](#)

## Fibre Channel zoning-based access control [4-28](#)

## FPSF

- load balancing (example) [2-5](#)

## frames

- configuring MTU size [6-5, 7-3](#)

# G

## Gigabit Ethernet

- IPv4 example configuration [6-4](#)

## Gigabit Ethernet interface example [4-54](#)

## Gigabit Ethernet interfaces

- configuring [6-4 to 6-10](#)
- configuring auto-negotiation [6-5, 7-3](#)
- configuring high availability [6-8 to 6-10](#)
- configuring IPv4 [7-2](#)
- configuring IPv6 addresses [8-12](#)
- configuring MTU frame sizes [6-5, 7-3](#)
- configuring promiscuous mode [6-6, 7-4](#)
- configuring VRRP [6-9](#)
- default parameters [7-6](#)
- IPv4-ACL guidelines [7-6](#)
- subinterfaces [6-6, 7-5](#)
- subnet requirements [6-6, 7-5](#)

## Gigabit Ethernet subinterfaces

- configuring VLANs [7-5](#)

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

## H

HA solution example [4-53](#)

HBA port [4-16, 4-21](#)

high availability

    Ethernet PortChannel [4-56](#)

    Ethernet PortChannels [2-6](#)

    Fibre Channel PortChannels [2-7](#)

    VRRP [2-6, 4-55](#)

    VRRPVRRP-based high availability [4-55](#)

## I

ICMP

    IPv6 [8-6](#)

ICMP packets

    IPv6 header format, figure [8-6](#)

in-band management

    IPFC [5-7](#)

initiators

    statically mapped iSCSI [4-37](#)

Internet Control Message Protocol. See ICMP

Internet Storage Name Service. See iSNS

IP connections

    active mode [2-22](#)

    initiating [2-22](#)

    passive mode [2-22](#)

IPFC

    configuration guidelines [5-7](#)

    description [5-7](#)

IPsec

    configuring with FCIP Wizard (procedure) [2-9](#)

IPS modules

    CDP support [6-10](#)

    FCIP [2-2](#)

    port modes [6-4, 7-1](#)

    software upgrades [6-3](#)

    supported features [6-1](#)

IPS port mode

    description [6-4](#)

IPS ports [4-9](#)

    modes [7-1](#)

    multiple connections [4-54](#)

IP storage services

    default parameters [6-10](#)

IP Storage services modules. See IPS modules

IPv4

    configuring Gigabit Ethernet interfaces [7-2](#)

    default settings [7-6](#)

    description [7-1](#)

    transitioning to IPv6 [8-15](#)

IPv4-ACLs

    guidelines for Gigabit Ethernet interfaces [7-6](#)

IPv4 addresses

    configuring IPv6 and IPV6 protocol stacks [8-13](#)

    IPv6 protocol stacks [8-10](#)

IPv4 default gateways

    configuring [5-3](#)

    description [5-3](#)

    static routes (tip) [5-6](#)

IPv4 default networks

    description [5-6](#)

IPv6

    address types [8-3](#)

    configuring addressing [8-11](#)

    configuring IPv4 and IPv6 addresses [8-13](#)

    configuring management interfaces [5-3](#)

    default settings [8-15](#)

    description [8-1 to 8-11](#)

    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](#)

    enhancements over IPv4 [8-1](#)

    ICMP [8-6](#)

    IPv6-ACL guidelines [8-14](#)

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

- neighbor discovery [8-7](#)
- path MTU discovery [8-7](#)
- router advertisement messages [8-9](#)
- router discovery [8-9](#)
- stateless autoconfiguration [8-9](#)
- static routes [8-13](#)
- transitioning from IPv4 [8-15](#)
- IPv6-ACLs
  - guidelines for IPv6 [8-14](#)
- IPv6 addresses
  - configuring [8-11](#)
  - configuring IPv4 and IPV6 protocol stacks [8-13](#)
  - formats [8-2](#)
  - link-local type [8-4](#)
  - multicast type [8-5](#)
  - prefix format [8-3](#)
  - unicast type [8-3](#)
- IPv6 neighbor discovery
  - advertisement messages [8-7](#)
  - description [8-7](#)
  - neighbor solicitation message, figure [8-8](#)
  - solicitation messages [8-7](#)
- IPv6 routing
  - enabling [8-11](#)
- IPv6 static routes
  - configuring [8-13](#)
- IQN
  - formats [4-9](#)
- IQNs
  - formats [4-9](#)
- ISCSI
  - enforcing access control [4-28](#)
- iSCSI
  - access control [4-25 to 4-29](#)
  - add initiator to zone database [4-26](#)
  - advanced VSAN membershipadvanced VSAN membership [4-25](#)
  - checking for WWN conflicts [4-20](#)
  - compatible drivers [4-2](#)
  - configuring [4-1, 4-1, 4-4, 4-57](#)
  - configuring AAA authentication [4-29, 4-30](#)
  - configuring ACLs [4-27](#)
  - configuring VRRP [4-55](#)
  - creating virtual targets [4-11](#)
  - default parameters [4-76](#)
  - discovery phase [4-28](#)
  - drivers [4-2](#)
  - enabling [4-4](#)
  - error [4-16](#)
  - Fibre Channel targets [4-8 to 4-14](#)
  - Gigabit Ethernet ports [6-4, 7-1](#)
  - GW flagiSCSI
    - gateway device [4-17](#)
  - HA with host without multi-path software [4-52](#)
  - initiator idle timeoutinitiator idle timeout
    - iSCSIinitiator idle timeout
      - configuring with Fabric Manager [4-17](#)
  - initiator name [4-31](#)
  - initiator targets [4-7](#)
  - IPS module support [6-2](#)
  - IQNs [4-15](#)
  - login redirect [4-38](#)
  - LUN mapping for targets [4-63](#)
  - MPS-14/2 module support [6-2](#)
  - multiple IPS ports [4-54](#)
  - PortChannel-based high availability [4-56](#)
  - PortChannel-based high availabilityEthernet
    - PortChannel-based high availability [4-56](#)
  - protocol [4-2](#)
  - requests and responses [4-3](#)
  - restrict an initiator to a specific user name for CHAP authentication [4-31](#)
  - routing [4-2](#)
  - routing modes chartrouting modes chart for iSCSI [4-35](#)
  - session creation [4-28](#)
  - session limits [4-16](#)
  - statically mapped initiators [4-37](#)
  - tables in Fabric Manager [4-20](#)

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

- targets in Device Manager [4-10](#)
- transparent initiator mode [4-16](#)
- transparent mode initiator [4-59](#)
- users with local authentication [4-31](#)
- using iSCSI Wizard (procedure) [4-6 to 4-8](#)
- VSAN membership [4-23](#)
- VSAN membership example [4-24](#)
- VSAN membership for iSCSI interfaces [4-23, 4-24](#)
- zone name [4-8](#)
- iSCSI authentication
  - configuring [4-29, 4-44](#)
  - configuring RADIUS (procedure) [4-32](#)
  - external RADIUS servers [4-58](#)
  - global override [4-30](#)
  - local authentication [4-31](#)
  - mechanisms [4-30](#)
  - restricting on initiators [4-31](#)
  - scenarios [4-57](#)
  - setup guidelines [4-57](#)
- iSCSI-based access control [4-27](#)
- iSCSI devices
  - example membership in VSANs [4-24](#)
- iscsi-gw [4-21](#)
- iSCSI high availability
  - configuring [4-51 to 4-57](#)
- ISCSI hosts
  - VSAN membership [4-23](#)
- iSCSI hosts
  - initiator identification [4-15](#)
  - initiator presentation modes [4-15](#)
  - initiator presentation modesinitiator presentation modes [4-15](#)
- iSCSI initiators
  - configuring static IP address mapping [4-19](#)
  - dynamic mapping [4-18](#)
  - idle timeout [4-17](#)
  - making dynamic WWN mapping static [4-20](#)
  - proxy mode [4-21](#)
  - statically mapped (procedure) [4-19](#)
  - static mapping [4-18](#)
  - transparent mode [4-16](#)
  - WWN assignments [4-18](#)
- iSCSI interfaces
  - configuring [4-15, 4-15](#)
  - configuring listener ports [4-33](#)
  - configuring listener portsiSCSI
    - listener port [4-33](#)
  - configuring QoS [4-33](#)
  - configuring routing mode [4-34](#)
  - configuring routing modesiSCS
    - configuring routing modesrouting modes [4-34](#)
  - configuring TCP tuning parameters [4-33](#)
  - creating [4-6](#)
  - creatingiSCSI
    - creating interfaces [4-6](#)
  - VSAN membership [4-24](#)
- iSCSI LUs [4-9](#)
- iSCSI protocol [4-1](#)
- iSCSI server load balancing [4-36](#)
- iSCSI Server Load Balancing. See iSLB
- iSCSI sessions
  - authenticationiSCSI
    - session authenticationauthentication
  - iSCSI session [4-29](#)
- iSCSI targets
  - advertising [4-13](#)
  - dynamic importing [4-9](#)
  - dynamic mapping [4-9](#)
  - examples [4-13](#)
  - secondary access [4-53](#)
  - static importing [4-11](#)
  - static importingstatic mappingiSCSI targets
    - static mapping [4-11](#)
  - transparent failover [4-51](#)
- iSLB
  - activating zones [4-42, 4-43](#)
  - auto-zoning [4-47](#)



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

- committing configuration changes
    - committing configuration changes
      - iSLB [4-49](#)
  - configuration distribution [4-47](#), [4-48](#)
  - configuration limits [4-37](#)
  - configuration prerequisites [4-38](#)
  - configuring [4-36](#)
  - configuring initiators and targets [4-42](#)
  - configuring VRRP [4-46](#)
  - configuring with Device Manager [4-38](#)
  - configuring zones [4-42](#), [4-43](#)
  - default settings [4-77](#)
  - distributing configuration using CF [4-47](#)
  - dynamic initiator mapping [4-41](#)
  - enabling configuration distribution [4-48](#)
  - initiator WWN assignment [4-36](#)
  - load balancing algorithm [4-46](#)
  - maximum initiators [4-37](#)
  - static initiator configuration
    - initiator configuration
      - static iSLB [4-37](#)
  - VSAN membership [4-41](#)
  - zone set activation failed [4-43](#)
- iSLb
- default settings [4-77](#)
- iSLB auto-zone feature [4-37](#)
- iSLB initiators [4-38](#)
- activating zones [4-43](#)
  - assigning WWNs [4-41](#)
  - configuring [4-40](#) to [4-44](#)
  - configuring load balancing metrics [4-42](#)
  - configuring zones [4-43](#)
  - dynamic initiator mapping [4-41](#)
  - VSAN membership [4-41](#)
- iSLB initiator targets
- activating zones [4-43](#)
  - configuring zones [4-43](#)
  - description [4-42](#)
- iSLB sessions
- authentication [4-44](#)
- authenticationiSLB
- sessions authentication [4-44](#)
- maximum per IPS portiSLB
- maximum sessions per IPS port [4-37](#)
- iSLB with CFS distribution [4-37](#)
- iSMS servers
- enabling [4-72](#)
- iSNS
- client registration [4-73](#)
  - cloud discovery [4-74](#)
  - configuring [4-74](#)
  - configuring servers [4-72](#) to [4-74](#)
  - description [4-68](#)
  - ESI [4-73](#)
- iSNS cloud discovery
- automatic [4-76](#)
  - CFS distribution [4-76](#)
  - description [4-74](#)
  - enabling [4-75](#)
  - initiating on-demand [4-75](#)
- iSNS profiles
- creating [4-69](#)
- iSNS servers
- configuration distribution [4-72](#)
  - configuring ESI retry count [4-73](#)
  - enabling [4-72](#)
  - example scenario [4-71](#)

---

## J

- jitter
  - configuring estimated maximum in FCIP profiles [2-20](#)
- jumbo frames. See MTUs

---

## K

- keepalive timeouts
  - configuring in FCIP profiles [2-18](#)

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

## L

latency  
    forwarding [4-34](#)

link-local addresses  
    description [8-4](#)  
    format, figure [8-4](#)

link redundancy  
    Ethernet PortChannel aggregation [6-9](#)

load balancing [4-36, 4-38](#)  
    FSPF (example) [2-5](#)  
    PortChannels (example) [2-5](#)  
    weighted [4-42](#)

load metric [4-42](#)

lock the fabric [4-49](#)

LUN [4-9](#)  
    trespass for storage port failover [4-54](#)

LUN mapping [4-53](#)  
    iSCSI [4-63](#)

LUNs  
    explicit access control [4-21](#)  
    mapping and assignment [4-21](#)

LUs [4-9](#)

## M

management interfaces  
    configuring [5-3](#)  
    configuring for IPv6 [5-3](#)

maximum retransmissions  
    configuring in FCIP profiles [2-18](#)

MD5 authentication  
    VRRP [5-14](#)

merge status conflictsiSLB  
    merge status conflictsCFS  
        merge status conflicts [4-50](#)

mgmt0 interfaces  
    local IPv4 routing [5-6](#)

minimum retransmit timeouts

    configuring in FCIP profiles [2-18](#)

MPS-14/2 modules [4-1, 4-2, 4-3, 4-6, 4-21, 4-28](#)  
    CDP support [6-10](#)  
    FCIP [2-2](#)  
    port modes [6-4, 7-1](#)  
    software upgrades [6-3](#)  
    supported features [6-1](#)

MTU frame sizes  
    configuring Gigabit Ethernet interfaces [6-5](#)

MTUs  
    configuring frame sizes [7-3](#)  
    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 [4-52](#)

multiple VSANs  
    configuring [5-9](#)

Multiprotocol Services modules. See MPS-14/2 modules

mutual CHAP authentication  
    configuring for iSCSI [4-31](#)  
    configuring for iSLB [4-44](#)  
    configuring for iSLBI [4-44](#)

## N

None authentication [4-29](#)

NTP  
    time-stamp option [2-22](#)

## O

overlay VSANs  
    configuring [5-8](#)  
    description [5-8](#)

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

## P

packets

- discarding in FCIP [2-22](#)

pass-thru routing mode [4-34, 4-36](#)

path MTUs. See PMTUs

PDU [4-34](#)

PMTUs

- configuring in FCIP profiles [2-18](#)

PortChannel

- interfaces [4-13](#)
- subinterfaces [4-13](#)

PortChannels

- configuring for FCIP high availability [2-5](#)
- IQN formats [4-9](#)
- load balancing (example) [2-5](#)
- member combinations [6-9](#)
- redundancy [2-7](#)

port modes

- IPS [6-4, 7-1](#)

ports

- virtual E [2-2](#)

promiscuous mode

- configuring Gigabit Ethernet interfaces [6-6, 7-4](#)

protocol [4-1](#)

protocols

- VRRP [4-9](#)

proxy initiator

- configuring iSCSI
- configuring proxy initiator [4-22](#)

proxy initiator mode [4-15, 4-26](#)

- configuring [4-21](#)
- zoning [4-23](#)

proxy initiator mode iSCSI

- proxy initiator mode [4-21](#)

pWWNs

- converting dynamic to static [4-20](#)

## Q

QoS

- DSCP value [2-25](#)

QoS values

- configuring [4-33](#)

## R

RADIUS [4-58](#)

- AAA authentication [4-29, 4-44](#)
- configuring an iSCSI RADIUS server iSCSI
- configuring a RADIUS server [4-32](#)

redundancy

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

router discovery

- IPv6 [8-9](#)

RSCNs [4-17](#)

## S

SACKs

- configuring in FCIP profiles [2-18](#)

SAN extension tuner

- configuring [3-2](#)
- data patterns [3-3](#)
- default settings [3-6](#)
- description [3-1](#)
- license requirements [3-3](#)
- tuning guidelines [3-2](#)

SCSI

- routing requests [4-2](#)

security parameter index. See SPI

selective acknowledgments. See SACKs

SPI

- configuring virtual routers [5-14](#)
- statically imported iSCSI targets [4-53](#)

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

static iSLB initiator  
     converting [4-41](#)  
 static mapped iSCSI target/iSCSI  
     static mapped target [4-28](#)  
 static mapping [4-41](#)  
 static WWN mapping [4-26](#)  
 store-and-forward routing mode [4-34, 4-36](#)  
 subnets  
     requirements [6-6, 7-5](#)  
 switch management  
     in-band [5-7](#)  
 switchovers  
     VRRP [2-6](#)

## T

TACACS+  
     AAA authentication [4-44](#)  
 target discovery [4-73](#)  
 TCP connections  
     FCIP profiles [2-4](#)  
 TCP parameters  
     configuring in FCIP profiles [2-17](#)  
 TCP tuning parameters [4-33](#)  
 transient failure [4-17](#)  
 transparent initiator mode [4-15](#)  
 transparent initiator mode/iSCSI  
     transparent initiator mode [4-21](#)  
 troubleshooting  
     CTC [2-17](#)  
 trunking mode  
     FCIP interface [2-4](#)

## 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 router IDs. See VR IDs  
 Virtual Router Redundancy Protocol. See VRRP  
 Virtual Router Redundancy Protocol/protocols  
     Virtual Router Redundancy [4-36](#)  
 virtual routers  
     adding [5-12](#)  
     adding primary IP addresses [5-13](#)  
     authentication [5-14](#)  
     default settings [5-15](#)  
     deleting [5-12](#)  
     initiating [5-13](#)  
     setting priorities [5-13](#)  
 VLANs  
     configuring on Gigabit Ethernet subinterfaces [7-5](#)  
     description [6-6, 7-4](#)  
 VR IDs  
     description [5-11](#)  
     mapping [5-11](#)  
 VRRP [4-36](#)  
     algorithm for selecting Gigabit Ethernet  
     interfaces [4-46](#)  
     backup switches [5-11](#)  
     configuring advertisement time intervals [5-13](#)  
     configuring for Gigabit Ethernet interfaces [6-9](#)  
     configuring for iSLB [4-46](#)  
     configuring virtual routers [5-12](#)  
     default settings [5-15](#)  
     description [5-11, 6-8](#)  
     group members [6-8](#)  
     initiating virtual routers [5-13](#)  
     IQN formats [4-9](#)  
     iSCSI parameter change impact [4-46](#)  
     iSLB [4-44](#)  
     master switches [5-11](#)

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

- MD5 authentication [5-14](#)
- primary IP address [5-13](#)
- priority preemption [5-13](#)
- security authentication [5-14](#)
- setting priorities [5-13](#)
- setting priority [5-13](#)
- simple text authentication [5-14](#)
- VRRP group [4-24](#)
- VRRP–I f iSCSI login redirect [4-38](#)
- VSAN membership
  - iSCSI hosts [4-23](#)
  - iSCSI hostsiSCSI
    - VSAN membership for hosts [4-23](#)
  - iSCSI interfaces [4-24](#)
- VSANs
  - configuring multiple IPv4 subnets [5-9](#)
  - example membership for iSCSI devices [4-24](#)
  - gateway switches [5-6](#)
  - IPv4 static routing [5-7](#)
  - iSLB [4-41](#)
  - iSLB initiators [4-41](#)
  - overlaid routes [5-7](#)
  - traffic routing between [5-1](#)
  - VRRP [5-11](#)
- configuring for iSCSI [4-26](#)
- configuring for iSCSIiSCSI
  - configuring zoning based access control [4-26](#)

---

## W

- window management
  - configuring in FCIP profiles [2-19](#)
- WWNs
  - static binding [4-21](#)

---

## Z

- zones
  - configuring and activating for iSLB [4-42](#)
  - iSLB [4-42, 4-43](#)
- zoning based access control

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