



Send comments to mdsfeedback-doc@cisco.com.



Cisco MDS 9000 Family Troubleshooting Guide

Cisco MDS SAN-OS Release 1.2(1a)
August 2003

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number:
Text Part Number: OL-3450-02



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.

Copyright © 2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Net Readiness Scorecard, Networking Academy, and ScriptShare are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0303R)

Cisco MDS 9000 Family Troubleshooting Guide

Copyright © 2003, Cisco Systems, Inc.

All rights reserved.

Send comments to mdsfeedback-doc@cisco.com.

CONTENTS

Preface vii

Document Organization	vii
Document Conventions	vii
Related Documentation	viii
Obtaining Documentation	viii
World Wide Web	ix
Documentation CD-ROM	ix
Ordering Documentation	ix
Documentation Feedback	ix
Obtaining Technical Assistance	x
Cisco.com	x
Technical Assistance Center	x
Cisco TAC Web Site	xi
Cisco TAC Escalation Center	xi
Obtaining Documentation	xi
World Wide Web	xi
Documentation CD-ROM	xii
Ordering Documentation	xii
Documentation Feedback	xii
Obtaining Technical Assistance	xii
Cisco.com	xiii
Technical Assistance Center	xiii
Cisco TAC Web Site	xiii
Cisco TAC Escalation Center	xiv

CHAPTER 1

Troubleshooting Overview 1-1

Introduction	1-1
Basics	1-2
Basic Connectivity	1-2
Fibre Channel End-to-End Connectivity	1-2
Fabric Issues	1-3
Using Host Diagnostic Tools	1-3
Using Cisco MDS 9000 Family Tools	1-4
Command-Line-Interface (CLI)	1-4

Send comments to mdsfeedback-doc@cisco.com.

CLI Debug	1-4
FC Ping and FC Traceroute	1-7
Cisco Fabric Manager	1-8
Fabric Manager and Device Manager	1-8
Analyzing Switch Device Health	1-9
Analyzing End-to-End Connectivity	1-10
Analyzing Switch Fabric Configuration	1-10
Analyzing the Results of Merging Zones	1-11
Alerts and Alarms	1-12
SCSI Target Discovery	1-12
SNMP and RMON Support	1-13
Using RADIUS	1-14
Using Syslog	1-15
Using Fibre Channel SPAN	1-15
Using Cisco Network Management Products	1-16
Cisco MDS 9000 Port Analyzer Adapter	1-17
Cisco Fabric Analyzer	1-17
CiscoWorks RME	1-19
Using Other Troubleshooting Products	1-20
Fibre Channel Testers	1-20
Fibre Channel Protocol Analyzers	1-21

CHAPTER 2

Troubleshooting Switch System Issues	2-1
Recovering the Administrator Password	2-1
Troubleshooting System Restarts	2-1
Overview	2-1
Working with Recoverable Restarts	2-2
Working with Unrecoverable System Restarts	2-5

CHAPTER 3

Troubleshooting Switch Level Issues and Interswitch Connectivity	3-1
Troubleshooting E Port Connectivity - ISL Isolation	3-1
Overview	3-2
Troubleshooting steps	3-3
Troubleshoot Switch and Port Parameters	3-3
Troubleshooting a ZONE Merge Failures	3-5
Troubleshooting a VSAN Configuration Conflict	3-6
Troubleshooting a Domain ID Conflict	3-7
Troubleshooting TE Port Connectivity - VSAN Isolation	3-10
Troubleshooting Fx Port Connectivity	3-13

Send comments to mdsfeedback-doc@cisco.com.

Fx Port Fails to Achieve Up State	3-13
FCOT Is Not Present	3-13
Link_Failure or Not Connected	3-14
Interface Bouncing between Offline and Initializing	3-15
Point-to-point link comes up as FL_Port	3-17
Interface UP and Connectivity Problems - Troubleshooting VSANs and Zones	3-18
Troubleshooting Zones – Case of end devices belonging to the default zone	3-19
Troubleshooting Zones – Case of end devices belonging to a specific zone	3-19
Verify active zoneset configuration	3-19
Verify active zoneset membership	3-20
Other useful commands	3-21
Using the GUI to Troubleshoot Zoning Configuration Issues	3-22
Troubleshooting VSANs	3-23
Using the GUI to Troubleshoot VSAN Membership Problems	3-24

CHAPTER 4

Troubleshooting Switch Fabric Level Issues 4-1

Troubleshooting Name Server Issues	4-1
Overview	4-1
Nx Port Registration Problems	4-2
Troubleshooting FSPF Issues	4-6
Overview	4-6
Loss of Two-way Communication	4-7
Wrong Hello Interval on an ISL Interface	4-7
Resolving the Wrong Hello Interval Problem	4-8
Wrong Dead Interval on an ISL Interface	4-8
Resolving a Wrong Dead Interval Problem	4-9
Region Mismatch on Switch	4-9
Resolving a Region Mismatch Problem	4-10
FSPF Issues in a Single-VSAN Environment	4-11
FSPF Issues in a Multi-VSAN Environment	4-13
Troubleshooting Zoning Issues	4-14
Mismatched Active Zonesets Within the Same VSAN	4-14
Importing or Exporting a Zoneset Between Switches	4-16
Deactivating a Zoneset and Restarting the Zone Merge Process	4-17
Misconfigured Zones Within an Active Zoneset in the Same VSAN	4-19

CHAPTER 5

Troubleshooting IP Storage Issues 21

Overview	21
Troubleshooting IP Connections	22

Send comments to mdsfeedback-doc@cisco.com.

Verifying Basic Connectivity	22
Verifying Static IP Routing	24
Troubleshooting FCIP Connections	25
One-to-One FCIP Tunnel Creation and Monitoring	25
One to three FCIP tunnel creation and monitoring	35
FCIP Profile Misconfiguration Examples	36
Interface FCIP Misconfiguration Examples	39
FCIP Special Frame Tunnel Creation and Monitoring	46
Special Frame Misconfiguration Examples	48
Troubleshooting iSCSI Issues	51
Troubleshooting iSCSI Authentication	51
Configuring Authentication	52
Troubleshooting Username/Password Configuration	53
Troubleshooting Radius Configuration	53
Troubleshooting Radius Routing Configuration	56
Troubleshooting Dynamic iSCSI Configuration	56
Checking the Configuration	56
Performing Basic Dynamic iSCSI Troubleshooting	57
Useful show Commands for Debugging Dynamic iSCSI Configuration	57
Virtual Target Access Control	59
Useful show Commands for Debugging Static iSCSI Configuration	59
Fine Tuning/Troubleshooting IPS iSCSI TCP Performance	64
Lab Setup	65
Configuration from the Bottom MDS	65
Changing TCP Parameters	69

Send comments to mdsfeedback-doc@cisco.com.

Preface

This document is intended to provide guidance for troubleshooting issues that may appear when deploying a storage area network (SAN) using the Cisco MDS 9000 Family of switches. This document will help you investigate the configuration of the different systems included in the SAN environment, such as hubs, hosts, and storage arrays and Cisco MDS 9000 switches. It also covers basic storage commands, switch configurations, and common storage parameters. It introduces tools and methodologies to recognize a problem, determine its cause, and find possible solutions.

Document Organization

This document is organized into the following chapters:

Chapter	Title	Description
Chapter 1	Troubleshooting Overview	Describes basic concepts, methodology, and tools to use for troubleshooting.
Chapter 2	Troubleshooting Switch Hardware and Booting Problems	Describes how to identify and resolve problems for a single Cisco MDS 9000 Family switch.
Chapter 3	Troubleshooting Switch Level Issues and Interswitch Connectivity	Describes how to identify and resolve problems that affect basic connectivity between switches, hosts, and storage in the network fabric.
Chapter 4	Troubleshooting Switch Fabric Level Issues	Describes switch fabric-level troubleshooting procedures.
Chapter 5	Troubleshooting IP Storage Issues	Describes IP storage troubleshooting procedures for the FCIP and iSCSI features.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
[]	Elements in square brackets are optional.

Send comments to mdsfeedback-doc@cisco.com.

[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use these conventions:

screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

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



Caution

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

Related Documentation

Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family

Quick Start Guide for the Cisco MDS 9000 Family

Cisco MDS 9200 Series Hardware Installation Guide

Cisco MDS 9500 Series Hardware Installation Guide

Cisco MDS 9000 Family Command Reference

Cisco MDS 9000 Family Fabric Manager User Guide

Cisco MDS 9000 Family Troubleshooting Guide

Cisco MDS 9000 Family System Messages Guide

Cisco MDS 9000 Family MIB Reference Guide

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

Send comments to mdsfeedback-doc@cisco.com.

World Wide Web

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

<http://www.cisco.com>

Documentation CD-ROM

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

Ordering Documentation

Cisco documentation is available in the following ways:

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

http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

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

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on Cisco.com, you can submit technical comments electronically. Click the **Fax** or **Email** option under the “Leave Feedback” at the bottom of the Cisco Documentation home page.

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

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

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

We appreciate your comments.

Send comments to mdsfeedback-doc@cisco.com.

Obtaining Technical Assistance

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

Cisco.com

Cisco.com is the foundation of a suite of interactive, network services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

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

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

<http://www.cisco.com>

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4) —You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3) —Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2) —Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1) —Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Send comments to mdsfeedback-doc@cisco.com.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

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

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

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

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

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

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a case is automatically opened.

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

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

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

Obtaining Documentation

These sections explain how to obtain documentation from Cisco Systems.

World Wide Web

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

<http://www.cisco.com>

Translated documentation is available at this URL:

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

Send comments to mdsfeedback-doc@cisco.com.

Documentation CD-ROM

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

Ordering Documentation

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:
http://www.cisco.com/cgi-bin/order/order_root.pl
- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:
<http://www.cisco.com/go/subscription>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. In the Cisco Documentation home page, click the **Fax** or **Email** option in the “Leave Feedback” section at the bottom of the page.

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

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

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

We appreciate your comments.

Obtaining Technical Assistance

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

Send comments to mdsfeedback-doc@cisco.com.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you with these tasks:

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

If you want to obtain customized information and service, you can self-register on Cisco.com. To access Cisco.com, go to this URL:

<http://www.cisco.com>

Technical Assistance Center

The Cisco Technical Assistance Center (TAC) is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two levels of support are available: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Cisco TAC inquiries are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

The Cisco TAC resource that you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

You can use the Cisco TAC Web Site to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to this URL:

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

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

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

Send comments to mdsfeedback-doc@cisco.com.

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

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

If you have Internet access, we recommend that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

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

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

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

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



Troubleshooting Overview

This chapter introduces the basic concepts, methodology, and tools to use for troubleshooting problems that may occur when configuring and using a Cisco MDS 9000 Family switch. The two most common symptoms of problems occurring in a storage network are:

- A host not accessing its allocated storage
- An application not responding after attempting to access the allocated storage

To identify the possible problems, you need to use a variety of tools and understand the overall storage environment. For this reason, this chapter describes a number of general troubleshooting tools and procedures in addition to those that are specific to the Cisco MDS 9000 family. This chapter also provides a plan of attack for investigating storage issues. Refer to the other chapters in this book for detailed explanations of specific issues.

This chapter includes the following sections:

- Introduction, page 1-1
- Using Host Diagnostic Tools, page 1-3
- Using Cisco MDS 9000 Family Tools, page 1-4
- Using Cisco Network Management Products, page 1-16
- Using Other Troubleshooting Products, page 1-20

Introduction

Some basic questions should be answered before you go into too much detail about specific problems and solutions. A process of elimination can determine which network components have been subject to change and therefore may be the cause of your problems. The main steps you should follow to identify a problem in a SAN environment include:

1. Verify physical connectivity and registration to the fabric
2. Verify storage subsystem and server configuration
3. Verify end-to-end connectivity and fabric configuration

This section provides a series of questions that may be useful when troubleshooting a problem with a Cisco MDS 9000 family switch or connected devices. Use the answers to these questions to plan a course of action and to determine the scope of the problem. For example, if a host can only see some of the LUNs, (but not all of them) on an existing subsystem, then fabric-specific issues (FSPF, ISLs, FCNS) do not need to be investigated, as they are currently working correctly. The fabric components can therefore be eliminated from the problem.

Send comments to mdsfeedback-doc@cisco.com.

By answering the questions in the following subsections, you can determine the paths you need to follow and the components that you should investigate further. These questions are independent of host, switch or subsystem vendor.

Basics

- Is this a newly installed system or an existing installation? (It could be a new SAN, host or subsystem, or new LUNs exported to an existing host.)
- Has the host ever been able to see its storage?
- Does the host recognize any LUNs in the subsystem?
- Are you trying to solve an existing application problem (too slow, too high latency, excessively long response time) or did the problem show up recently?
- What changed in the configuration or in the overall infrastructure immediately before the applications started to have problems?

Basic Connectivity

- Are you using the correct fiber (SM or MM)?
- Did you check for a broken fiber?
- Is the LED on the connected module Fibre Channel port green, and do the LEDs on any HBA/Storage Subsystem ports indicate normal functionality?
- Is there a LUN masking policy applied on the storage subsystem? If yes, is the server allowed to see the LUNs exported by the storage array?
- Is there any LUN masking policy configured on the host? Did you enable the server to see all the LUNs it can access?
- If LUN masking software is used, is the host's PWWN listed in the LUN masking database?
- Is the subsystem configured for NPort?

Examine the FLOGI database on the two switches that are directly connected to the host HBA and subsystem ports. Also, verify that both ports (attached port on MDS-A and MDS-B) are members of the same VSAN. If both devices are listed in the FCNS database then ISLs are not an issue.

Fibre Channel End-to-End Connectivity

Answering the following questions will help to determine if end-to-end fibre channel connectivity exists from a host or subsystem perspective:

- Does the host list the subsystem's WWPN or FCID in its logs?
- Does the subsystem list the host's WWPN or FCID in its logs or LUN masking database?
- Can the host complete a port login (PLOGI) to the storage subsystem?
- Is there any SCSI exchange that takes place between the server and the disk array?
- Is the HBA configured for NPort?

You can use the HBA configuration utilities or the host system logs to determine if the subsystem PWWN or FCID is listed as a device. This can validate that FSPF is working correctly.

Send comments to mdsfeedback-doc@cisco.com.

Fabric Issues

- Are both the host bus adapter (HBA) and the subsystem port successfully registered with the fabric name server?
- Does the correct PWWN for the Server HBA and the storage subsystem port show up on the correct port in the FLOGI database? In other words, is the device plugged into the correct port?
- Does any single zone contain both devices? The zone members can be WWNs, FCIDs.
- Is the zone correctly configured and part of the active configuration or zoneset within the same VSAN?
- Do the ISLs show any VSAN isolation?
- Do the host and storage belong to the same VSAN?
- Are any parameters, such as FSPF, Static Domain Assignment, VSAN or Zoning, mismatched in the configuration of the different switches in the fabric?

Using Host Diagnostic Tools

Most host systems provide utilities or other tools that you can use for troubleshooting the connection to the allocated storage. For example, on a Windows system, you can use the Diskmon or Disk Management tool to verify accessibility of the storage and to perform some basic monitoring and administrative tasks on the visible volumes.

Alternatively, you can use Iometer, an I/O subsystem measurement and characterization tool, to generate a simulated load and measure performance. Iometer is a public domain software utility for Windows, originally written by Intel, that provides correlation functionality to assist with performance analysis.

Iometer measures the end-to-end performance of a SAN without cache hits. This can be an important measurement because if write or read requests go to the cache on the controller (a cache hit) rather than to the disk subsystems, performance metrics will be artificially high. You can obtain Iometer from SourceForge.net at the following URL:

<http://sourceforge.net/projects/iometer/>

Iometer is not the only I/O generator you can use to simulate traffic through the SAN fabric. Other popular I/O generators and benchmark tools used for SAN testing include Iozone and Postmark. Iozone is a file system benchmark tool that generates and measures a variety of file operations. It has been ported to many systems and is useful for performing a broad range of file system tests and analysis.

Postmark was designed to create a large pool of continually changing files, which simulates the transaction rates of a large Internet mail server.

PostMark generates an initial pool of random text files in a configurable range of sizes. Creation of the pool produces statistics on continuous small file creation performance. Once the pool is created, PostMark generates a specified number of transactions, each of which consists of a pair of smaller transactions:

- Create file or Delete file
- Read file or Append file

Benchmark is available from Network Appliance, Inc. at the following URL:

http://www.netapp.com/tech_library/3022.html

Benchmarking tools offer a variety of capabilities and you should select the one that provides the best I/O characteristics of your application environment.

Send comments to mdsfeedback-doc@cisco.com.

Utilities provided by the Sun Solaris operating system let you determine if the remote storage has been recognized and exported to you in form of a raw device or mounted file system, and to issue some basic queries and tests to the storage. You can measure performance and generate loads using the **iostat** utility, the **perfmer** GUI utility, the **dd** utility, or a third-party utility like Extreme SCSI.

Every UNIX version provides similar utilities, but this guide only provides examples for Solaris. Refer to the documentation for your specific operating system for details.

Using Cisco MDS 9000 Family Tools

If the server does not see its storage and you cannot use the information available on the host side to determine the root cause of the problem, you can obtain additional information from a different viewpoint using the troubleshooting tools provided with the Cisco MDS 9000 family of switches. This section introduces these tools and describes the kinds of problems for which you can use each tool. It includes the following topics:

- Command-Line-Interface (CLI), page 1-4
- CLI Debug, page 1-4
- FC Ping and FC Traceroute, page 1-7
- Cisco Fabric Manager, page 1-8
- SCSI Target Discovery, page 1-12
- SNMP and RMON Support, page 1-13
- Using RADIUS, page 1-14
- Using Syslog, page 1-15
- Using Fibre Channel SPAN, page 1-15

Command-Line-Interface (CLI)

The Cisco MDS 9000 Family CLI lets you configure and monitor a Cisco MDS 9000 Family switch using a local console or remotely using a Telnet or SSH session. The CLI provides a command structure similar to Cisco IOS[®] software, with context-sensitive help, show commands, multi-user support, and roles-based access control.

CLI Debug

The Cisco MDS 9000 Family of switches includes an extensive debugging feature set for actively troubleshooting a storage network. Using the CLI, you can enable debugging modes for each switch feature and view a real-time updated activity log of the control protocol exchanges. Each log entry is time-stamped and listed in chronological order. Access to the debug feature can be limited through the CLI roles mechanism and can be partitioned on a per-role basis. While debug commands show realtime information, the **show** commands can be used to list historical information as well as realtime.



Note

You can log debug messages to a special log file, which is more secure and easier to process than sending the debug output to the console.

Send comments to mdsfeedback-doc@cisco.com.

By using the '?' option, you can see the options that are available for any switch feature, such as FSPF. A log entry is created for each entered command in addition to the actual debug output. The debug output shows a time-stamped account of activity occurring between the local switch and other adjacent switches.

You can use the debug facility to keep track of events, internal messages and protocol errors. However, you should be careful with using the debug utility in a production environment, because some options may prevent access to the switch by generating too many messages to the console or if very CPU-intensive may seriously affect switch performance.



Note

It is a good idea to open a second Telnet or SSH session before entering any debug commands. That way, if the debug output comes too fast to stop it in the output window, you can use the second session to enter the **undebug all** command to stop the debug message output.

The following is an example of the output from the **debug flogi event** command

```
switch# debug flogi event interface fc1/1
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_FLOGI_RECEIVED]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_VALID_FLOGI]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_GET_FCID]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_GET_FCID]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_VALID_FCID]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_CONFIG_DONE_PENDING]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_RIB_RESPOSE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_NAME_SERVER_REG_RESPONSE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_ACL_CFG_RESPONSE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_ZS_CFG_RESPONSE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute_all: done processing event FLOGI_EV_ZS_CFG_RESPONSE
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_LCP_RESPONSE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:      current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:      current event [FLOGI_EV_CONFIG_DONE_COMPLETE]
Dec 10 23:40:26 flogi:      next state   [FLOGI_ST_FLOGI_DONE]
```

The following is a summary of the available debug commands:

Table 1-1 Debug commands

Debug command	Purpose
acl	Enable acl debugging
all	Enable all debugging

Send comments to mdsfeedback-doc@cisco.com.

Table 1-1 *Debug commands (continued)*

Debug command	Purpose
ascii-cfg	Configure ascii-cfg debugging
bootvar	Enable bootvar debugging
callhome	Enable debugging for Callhome
fc2	Configure FC2 debugging
fcc	Enable FCC debugging
fcdomain	Enable fcdomain debugging
fcfwd	Enable fcfwd debugging
fcns	Debug name server
fcs	Configure Fabric Configuration Server Debugging
flogi	Configure flogi debug
fspf	Configure FSPF debugging
hardware	Debug hardware, kernel loadable module parameters
ipconf	Enable IP configuration debugging
ipfc	Enable IPFC debugging
klm	Debug kernel loadable module parameters
logfile	Direct debug output to logfile
mcast	Enable mcast debugging
mip	Debug multiple IP kernel driver
module	Configure LC Manager debugging
ntp	Debug NTP module
platform	Configure Platform Manager debugging
port	Configure port debugging
port-channel	Enable port-channel debug
qos	Configure QOS Manager Debugging
rdl	Configure RDL debugging
rib	Configure rib debugging
rscn	Configure RSCN debugging
scsi-target	Configure scsi target daemon debugging
security	Configure debugging for security/accounting
sensor	Enable Sensor Mgr debugging
span	Configure SPAN debug
system	Enable System debugging
tlport	Configure TL Port debugging
vni	Enable virtual network interface debugging
vrrp	Enable vrrp debugging
vsan	Enable VSAN manager debugging

Send comments to mdsfeedback-doc@cisco.com.

Table 1-1 *Debug commands (continued)*

Debug command	Purpose
vsh	Enable vsh debugging
vshd	Configure vshd debugging
wwn	Configure WWN Manager Debugging
xbar	Enable xbar debugging
xbc	Enable Xbar Client debugging
zone	Zone server debug commands

FC Ping and FC Traceroute



Note

FC Ping and FC traceroute are used to troubleshoot problems with connectivity and path choices. They are not designed for use in identifying or resolving performance issues.

Ping and *Traceroute* are two of the most useful tools for troubleshooting TCP/IP networking problems. The *Ping* utility generates a series of *echo* packets to a destination across a TCP/IP internetwork. When the *echo* packets arrive at the destination, they are re-routed and sent back to the source. Using *Ping*, you can verify connectivity and latency to a particular destination across an IP routed network. *Traceroute* operates in a similar fashion, but can also determine the specific path that a frame takes to its destination on a hop-by-hop basis. These tools have been migrated to Fibre Channel for use with the Cisco MDS 9000 Family switches and are called *FC Ping* and *FC Traceroute*. You can use *FC Ping* and *FC Traceroute* from the CLI or from the Cisco Fabric Manager.

FC Ping allows you to ping a Fibre Channel *N_Port* or end device. By specifying the *FC_ID* or Fibre Channel address, you can send a series of frames to a target *N_Port*. Once these frames reach the outgoing *F_Port*, they are looped back to the source and a time-stamp is taken. *FC Ping* helps you to verify the connectivity and latency to an end *N_Port*. *FC Ping* uses the PRLI Extended Link Service, and verifies the presence of a FC entity in case of positive or negative answers.

FC Traceroute is slightly different than the IP equivalent because both the outbound and return paths are recorded as these may differ in a switched Fibre Channel network. The *FC Traceroute* command identifies the path taken on a hop-by-hop basis and includes a timestamp at each hop in both directions.

FC Ping and *FC Traceroute* are useful tools to check for network connectivity problems or verify the path taken toward a specific destination. You can use *FC Traceroute* to test the connectivity of TE ports along the path between the generating switch and the switch closest to the destination.



Note

FC Trace will only work across EISL links.

The following is an example of output from the **fcping** command:

```
switch# fcping fcid 0xef02c9 vsan 1
28 bytes from 0xef02c9 time = 1408 usec
28 bytes from 0xef02c9 time = 379 usec
28 bytes from 0xef02c9 time = 347 usec
28 bytes from 0xef02c9 time = 361 usec
28 bytes from 0xef02c9 time = 363 usec

5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 347/571/1408 usec
```

Send comments to mdsfeedback-doc@cisco.com.

The following is an example of output from the **fctrace** command:

```
switch# fctrace fcid 0xef0000 vsan 1
Route present for : 0xef0000
20:00:00:05:30:00:59:de(0xffffcee)
Latency: 0 msec
20:00:00:05:30:00:58:1e(0xffffc6c)
Timestamp Invalid.
20:00:00:05:30:00:59:1e(0xffffcef)
Latency: 0 msec
20:00:00:05:30:00:59:1e(0xffffcef)
Latency: 174860 msec
20:00:00:05:30:00:58:1e(0xffffc6c)
```

Cisco Fabric Manager

Cisco Fabric Manager provides fabric-wide management capabilities including discovery, multiple switch configuration, network monitoring, and troubleshooting. It provides the troubleshooting features described in the following topics:

- Fabric Manager and Device Manager, page 1-8
- Analyzing Switch Device Health, page 1-9
- Analyzing End-to-End Connectivity, page 1-10
- Analyzing Switch Fabric Configuration, page 1-10
- Analyzing the Results of Merging Zones, page 1-11
- Alerts and Alarms, page 1-12



Note

For detailed information about using Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

Fabric Manager and Device Manager

Fabric Manager provides a map of the discovered fabric and includes tables that display statistical information about the switches in the fabric. You can also select troubleshooting tools from the Fabric Manager Troubleshooting menu.



Note

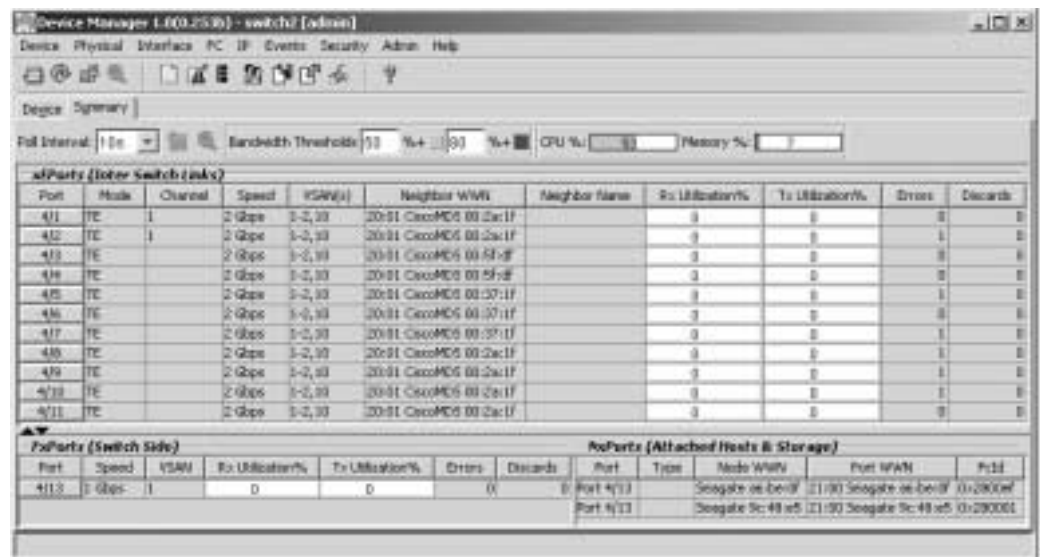
When you click on a zone or VSAN in Fabric Manager, the members of the zone or VSAN are highlighted on the Fabric Manager Map pane.

Device Manager provides a graphic display of a specific switch and shows the status of each port on the switch. From Device Manager, you can drill down to get detailed statistics about a specific switch or port.

Figure 1-1 shows the Fabric Manager Summary View window.

Send comments to mdsfeedback-doc@cisco.com.

Figure 1-1 Cisco Fabric Manager Summary View




The Summary View window lets you analyze switch performance issues, diagnose problems, and change parameters to resolve problems or inconsistencies. This view shows aggregated statistics for the active Supervisor Module and all switch ports. Information is presented in tabular or graphical formats, with bar, line, area, and pie chart options. You can also use the Summary View to capture the current state of information for export to a file or output to a printer.

Analyzing Switch Device Health

Choose the Switch Health option from the Fabric Manager Troubleshooting menu to determine the status of the components of a specific switch.

Figure 1-2 Switch Health Analysis Window



The Switch Health Analysis window displays any problems affecting the selected switches.

Send comments to mdsfeedback-doc@cisco.com.

Analyzing End-to-End Connectivity

Select the End to End Connectivity option from the Fabric Manager Troubleshooting menu to determine connectivity and routes among devices with the switch fabric. The connectivity tool checks to see that every pair of end devices in an active zone can talk to each other, using a Ping test and by determining if they are in the same VSAN. This option uses versions of the **ping** and **traceroute** commands modified for Fibre Channel networks.

The End to End Connectivity Analysis window displays the selected end points with the switch to which each is attached, and the source and target ports used to connect it.

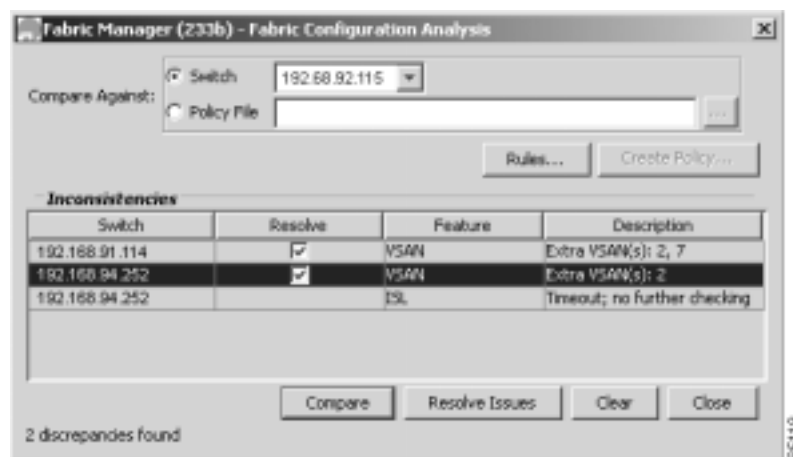
The output shows all the requests which have failed. The possible descriptions are:

- Ignoring empty zone—No requests are issued for this zone.
- Ignoring zone with single member—No requests are issued for this zone.
- Source/Target are unknown—No nameserver entries exist for the ports or we have not discovered the port during discovery.
- Both devices are on the same switch—The devices are not redundantly connected.
- No paths exist.
- Only one unique path exists.
- VSAN does not have an active zone set.
- Average time... micro secs—The latency value was more than the threshold supplied.

Analyzing Switch Fabric Configuration

Select the Fabric Configuration option from the Fabric Manager Troubleshooting menu to analyze the configuration of a switch by comparing the current configuration to a specific switch or to a policy file. You can save a switch configuration to a file and then compare all switches against the configuration in the file.

Figure 1-3 Fabric Configuration Analysis Window



You use a policy file to define the rules to be applied when running the Fabric Checker. When you create a policy file, the system saves the rules selected for the selected switch.

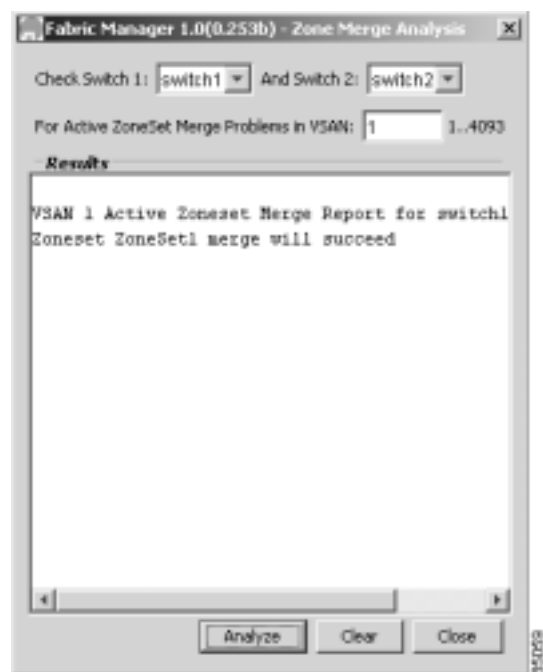
Send comments to mdsfeedback-doc@cisco.com.

Analyzing the Results of Merging Zones

Cisco Fabric Manager provides a very useful tool for troubleshooting problems that occur when merging zones configured on different switches.

Select the **Zone Merge** option on the Fabric Manager Troubleshooting menu to determine if two connected switches have compatible zone configurations.

Figure 1-4 Zone Merge Analysis Window



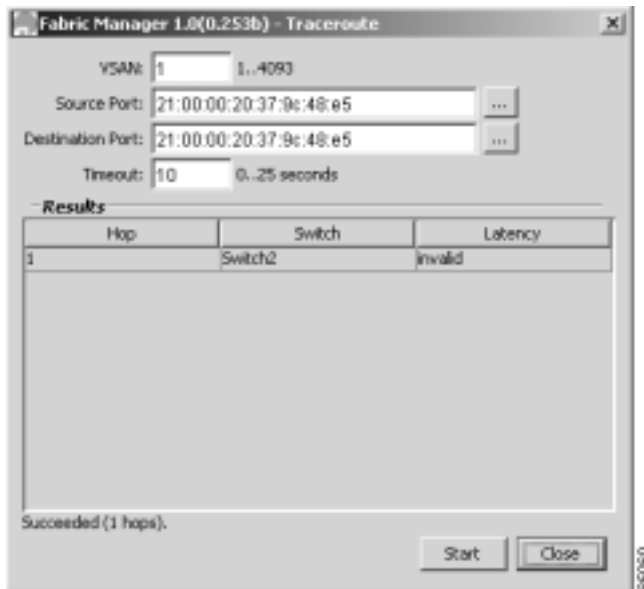
The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches.

You can use the following options on the Fabric Manager Troubleshooting menu to verify connectivity to a selected object or to open other management tools:

- Traceroute—Verify connectivity between two end devices that are currently selected on the Map pane.
- Device Manager—Launch Device Manager for the switch selected on the Map pane.
- Command Line Interface—Open a Telnet or SSH session for the switch selected on the Map pane.

Send comments to mdsfeedback-doc@cisco.com.

Figure 1-5 Traceroute Window



Alerts and Alarms

You can configure and monitor SNMP, RMON, Syslog, and Call Home alarms and notifications using the different options on the Device Manager Events menu. SNMP provides a set of preconfigured traps and informs that are automatically generated and sent to the destinations (trap receivers) that you identify. The RMON Threshold Manager lets you configure thresholds for specific events that trigger log entries or notifications. You can use either Fabric Manager or Device Manager to identify Syslog servers that will record different events or to configure Call Home, which can alert you through e-mail messages or paging when specific events occur.

SCSI Target Discovery

For more information about SCSI target discovery, refer to the *Cisco MDS 9000 Family Configuration Guide*.

The Fibre Channel name service is a distributed service in which all connected devices participate. As new SCSI target devices attach to the fabric, they register themselves with the name service, which is then distributed among all participating fabric switches. This information can then be used to help determine the identity and topology of nodes connected to the fabric.

For the Cisco MDS 9000 Family of switches, the SCSI Target Discovery feature has been added to provide added insight into connected SCSI targets. This feature allows the switch to briefly log into connected SCSI target devices and issue a series of SCSI inquiry commands to help discover additional information. The additional information that is queried includes logical unit number (LUN) details including the number of LUNs, the LUN IDs, and the sizes of the LUNs.

This information is then compiled and made available to through CLI commands, through the Cisco Fabric Manager, and also via an embedded SNMP MIB which allows the information to be easily retrieved by an upstream management application. Using the *SCSI Target Discovery* feature, you can have a much more detailed view of the fabric and its connected SCSI devices.

The following is an example of output from the `discover scsi-target` command:

Send comments to mdsfeedback-doc@cisco.com.

```
switch# discover scsi-target local remote
discovery started
switch# show scsi-target lun vsan 1
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b5 in VSAN 1, PWWN is 21:00:00:20:37:46:78:97
-----
LUN      Capacity  Status  Serial Number  Device-Id
      (MB)
-----
0x0      18210      Online  LRA2510000007027 C:1 A:0 T:3 20:00:00:20:37:46:78:97
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b6 in VSAN 1, PWWN is 21:00:00:20:37:5b:cf:b9
-----
LUN      Capacity  Status  Serial Number  Device-Id
      (MB)
-----
0x0      18210      Online  LR948730000007029 C:1 A:0 T:3 20:00:00:20:37:5b:cf:b9
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b9 in VSAN 1, PWWN is 21:00:00:20:37:18:6f:90
-----
LUN      Capacity  Status  Serial Number  Device-Id
      (MB)
-----
0x0      18210      Online  LR185918000001004 C:1 A:0 T:3 20:00:00:20:37:18:6f:90
```



Note

This tool can be effective to find out the number of LUNs exported by a storage subsystem, but it may be ineffective when LUN Zoning/LUN Security tools are used.

SNMP and RMON Support

The Cisco MDS 9000 Family of switches provides extensive SNMPv1, v2, and v3 support, including Management Information Bases (MIBs) and notifications (traps and informs).

The applications provided by Cisco that use SNMP include Fabric Manager, Cisco SSE and Cisco RME. Also, the SNMP standard allows any third-party applications that support the different MIBs to manage and monitor Cisco MDS 9000 Family switches.

SNMP v3 provides extended security. Each switch can be selectively enabled or disabled for SNMP service. In addition, each switch can be configured with a method of handling SNMPv1 and v2 requests.



Note

During initial configuration of your switch, the system prompts you to define SNMP v1 or V2 community strings and to create a SNMP v3 username and password.

Cisco MDS 9000 Family switches support over 50 different MIBs, which can be divided into the following six categories:

- IETF Standards-based Entity MIBs (for example, RFC273 ENTITY-MIB) These MIBs are used to report information on the physical devices themselves in terms of physical attributes etc.
- Cisco-Proprietary Entity MIBs (for example, CISCO-ENTITY-FRU-CONTROL-MIB) These MIBs are used to report additional physical device information about Cisco-only devices such as their configuration.

Send comments to mdsfeedback-doc@cisco.com.

- IETF IP Transport-oriented MIBs (for example, RFC2013 UDP-MIB) These MIBs are used to report transport-oriented statistics on such protocols as IP, TCP, and UDP. These transports are used in the management of the Cisco MDS 9000 Family through the OOB Ethernet interface on the Supervisor module.
- Cisco-Proprietary Storage and Storage Network MIBs (for example, NAME-SERVER-MIB) These MIBs were written by Cisco to help expose information that is discovered within a fabric to management applications not connected to the fabric itself. In addition to exposing configuration details for features like zoning and Virtual SANs (VSANs) via MIBs, discovered information from sources like the FC-GS-3 Name Server can be pulled via a MIB. Additionally, MIBs are provided to configure/enable features within the Cisco MDS 9000 Family. There are over 20 new MIBs provided by Cisco for this information and configuration capability.
- IETF IP Storage Working Group MIBs (for example, ISCSI-MIB) While many of these MIBs are still work-in-progress, Cisco is helping to draft such MIBs for protocols such as iSCSI and Fibre Channel-over-IP (FCIP) to be standardized within the IETF.
- Miscellaneous MIBs (for example, SNMP-FRAMEWORK-MIB) There are several other MIBs provided in the Cisco MDS 9000 Family switches for tasks such as defining the SNMP framework or creating SNMP partitioned views.

You can use SNMPv3 to assign different SNMP capabilities to specific roles.

Cisco MDS 9000 Family switches also support Remote Monitoring (RMON) for Fibre Channel. RMON provides a standard method to monitor the basic operations of network protocols providing connectivity between SNMP management stations and monitoring agents. RMON also provides a powerful alarm and event mechanism for setting thresholds and sending notifications based on changes in network behavior.

The RMON groups that have been adapted for use with Fibre Channel include the *AlarmGroup* and *EventGroup*. The *AlarmGroup* provides services to set alarms. Alarms can be set on one or multiple parameters within a device. For example, you can set an RMON alarm for a specific level of CPU utilization or crossbar utilization on a switch. The *EventGroup* lets you configure events that are actions to be taken based on an alarm condition. The types of events that are supported include *logging*, *SNMP traps*, and *log-and-trap*.

**Note**

To configure events within an RMON group, use the **Events > Threshold Manager** option from Device Manager.

Using RADIUS

RADIUS is fully supported for the Cisco MDS 9000 Family switches through the Fabric Manager and the CLI. RADIUS is a protocol used for the exchange of attributes or credentials between a head-end RADIUS server and a client device. These attributes relate to three classes of services:

- Authentication
- Authorization
- Accounting

Authentication refers to the authentication of users for access to a specific device. You can use RADIUS to manage user accounts for access to Cisco MDS 9000 Family switches. When you try to log into a switch, the switch validates you with information from a central RADIUS server.

Send comments to mdsfeedback-doc@cisco.com.

Authorization refers to the scope of access that you have once you have been authenticated. With Cisco MDS 9000 Family switches, assigned roles for users can be stored in a RADIUS server along with a list of actual devices that the user should have access to. Once the user has been authenticated, then switch can then refer to the RADIUS server to determine the extent of access the user will have within the switch network.

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

The following is an example of an accounting log entries.

```
switch# show accounting log
Sun Dec 15 04:02:27 2002:start:/dev/pts/0_1039924947:admin
Sun Dec 15 04:02:28 2002:stop:/dev/pts/0_1039924947:admin:vsh exited normally
Sun Dec 15 04:02:33 2002:start:/dev/pts/0_1039924953:admin
Sun Dec 15 04:02:34 2002:stop:/dev/pts/0_1039924953:admin:vsh exited normally
Sun Dec 15 05:02:08 2002:start:snmp_1039928528_172.22.95.167:public
Sun Dec 15 05:02:08 2002:update:snmp_1039928528_172.22.95.167:public:Switchname
set to Switch
```



Note

The accounting log only shows the beginning and ending (start and stop) for each session.

Using Syslog

Syslog lets you store a chronological log of system messages locally or sent to a central Syslog server. Syslog messages can also be sent to the console for immediate use. These messages can vary in detail depending on the configuration that you choose. Syslog messages are categorized into 7 severity levels from *debug* to *critical* events. You can limit the severity levels that are reported for specific services within the switch. For example, you may wish only to report *debug* events for the *FSPF* service but record all severity level events for the *Zoning* service.

A unique feature within the Cisco MDS 9000 Family switches is the ability to send RADIUS accounting records to the Syslog service. The advantage of this feature is that you can consolidate both types of messages for easier correlation. For example, when you log into a switch and change an FSPF parameter, Syslog and RADIUS provide complimentary information that will help you formulate a complete picture of the event.

Using Fibre Channel SPAN

For more information about configuring SPAN, refer to the *Cisco MDS 9000 Family Configuration Guide*.

You can use the Switched Port Analyzer (SPAN) utility to perform detailed troubleshooting or to take a sample of traffic from a particular application host for proactive monitoring and analysis. This utility is most helpful when you have a Fibre Channel protocol analyzer available and you are monitoring user traffic between two FC IDs.

When you have a problem in your storage network that you cannot solve by fixing the device configuration, you typically need to take a look at the protocol level. You can use debug commands to look at the control traffic between an end node and a switch. However, when you need to focus on all the traffic originating from or destined to a particular end node such as a host or a disk, you can use a protocol analyzer to capture protocol traces.

Send comments to mdsfeedback-doc@cisco.com.

To use a protocol analyzer, you must insert the analyzer in-line with the device under analysis, which disrupts input and output (I/O) to and from the device. This problem is worse when the point of analysis is on an Inter-Switch Link (ISL) link between two switches. In this case, the disruption may be significant depending on what devices are downstream from the severed ISL link.

In Ethernet networks, this problem can be solved using the SPAN utility, which is provided with the Cisco Catalyst Family of Ethernet switches. SPAN has also been implemented with the Cisco MDS 9000 Family of switches for use in Fibre Channel networks. SPAN lets you take a *copy* of all traffic and direct it to another port within the switch. The process is non-disruptive to any connected devices and is facilitated in hardware, which prevents any unnecessary CPU load. Using Fibre Channel SPAN, you can connect a Fibre Channel analyzer, such as a Finisar analyzer, to an unused port on the switch and then SPAN a copy of the traffic from a port under analysis to the analyzer in a non-disruptive fashion.

SPAN allows you to create up to 16 independent *SPAN* sessions within the switch. Each session can have up to four unique sources and one destination port. In addition, you can apply a filter to capture only the traffic received or the traffic transmitted. With Fibre Channel SPAN, you can even capture traffic from a particular Virtual SAN (VSAN).

To start the SPAN utility use the CLI command **span session session_num**, where *session_num* identifies a specific SPAN session. When you enter this command, the system displays a submenu, which lets you configure the destination interface and the source VSAN or interfaces.

```
switch2# config terminal
switch2(config)# span session 1
<<Create a span session>>

switch2(config-span)# source interface fc1/8
<<specify the port to be spanned>>

switch2(config-span)# destination interface fc1/3
<<specify the span destination (SD) port>>

switch2(config-span)# end
switch2# show span session 1
Session 1 (active)
  Destination is fc1/1
  No session filters configured
  Ingress (rx) sources are
    fc1/8,
  Egress (tx) sources are
    fc1/8,
```

Using Cisco Network Management Products

This section describes network management tools that are available from Cisco and are useful for troubleshooting problems with Cisco MDS 9000 Family switches and connected devices. It includes the following topics:

- Cisco MDS 9000 Port Analyzer Adapter, page 1-17
- Cisco Fabric Analyzer, page 1-17
- CiscoWorks RME, page 1-19

Send comments to mdsfeedback-doc@cisco.com.

Cisco MDS 9000 Port Analyzer Adapter

The Cisco MDS 9000 Port Analyzer Adapter is a stand-alone adapter card that converts Fibre Channel (FC) frames to Ethernet frames by encapsulating each Fibre Channel frame into an Ethernet frame. This product is meant to be used for analyzing SPAN traffic from a Fibre channel port on a Cisco MDS 9000 Family switch.

The Cisco Port Analyzer Adapter provides two physical interfaces:

- An FC interface that connects to the SPAN port of a Cisco MDS 9000 Family switch
- A 100/1000 Mb/s Ethernet port that forwards the encapsulated Fibre Channel traffic with a broadcast destination MAC Address



Note

The Cisco Port Analyzer Adapter does not support half-duplex mode and for this reason, it will not work when connected to a hub.

The Cisco Port Analyzer Adapter provides the following features:

- Encapsulates FC frames into Ethernet frames
- Sustains 32 max size FC frames burst (in 100 Mb/s mode)
- Line rate at 1Gb/s (for FC frames bigger than 91bytes)
- 64KBytes of onboard frame buffer
- Configurable option for Truncating FC frames to 256 Bytes - for greater burst
- Configurable option for Deep Truncating FC frames to 64 Bytes - best frames burst
- Configurable option for Ethernet Truncating FC frames to 1496 Bytes - max size E-net frames
- Configurable option for No Truncate Mode - sends jumbo frames on E-net side.
- Packet Counter (Indicates number of previous packet drops)
- SOF/EOF type information embedded
- 100/1000 Mb/s Ethernet interface - option on board
- Auto Configuration on power up
- Fibre Channel and Ethernet Link up indicator - LEDs.
- Checks FC frame CRC

When used in conjunction with the open source protocol analyzer, Ethereal (<http://www.ethereal.com>), the Cisco Port Analyzer Adapter provides a cost-effective and powerful troubleshooting tool. It allows any PC with a Ethernet card to provide the functionality of a flexible Fibre Channel analyzer. For more information on using the Cisco Port Analyzer Adapter see the *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Guide*.

Cisco Fabric Analyzer

For more information about using the Cisco Fabric Analyzer, refer to the *Cisco MDS 9000 Family Configuration Guide*.

The ultimate tool for troubleshooting network protocol problems is the protocol analyzer. Protocol analyzers promiscuously capture network traffic and completely decode the captured frames down to the protocol level. Using a protocol analyzer, you can conduct a detailed analysis by taking a sample of a

Send comments to mdsfeedback-doc@cisco.com.

storage network transaction and by mapping the transaction on a frame-by-frame basis, complete with timestamps. This kind of information lets you pinpoint a problem with a high degree of accuracy and arrive at a solution more quickly. However, dedicated protocol analyzers are expensive and they must be placed locally at the point of analysis within the network.

With the Cisco Fabric Analyzer, Cisco has brought Fibre Channel protocol analysis within a storage network to a new level of capability. Using Cisco Fabric Analyzer, you can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be present locally at the point of analysis.

The Cisco Fabric Analyzer consists of three main components:

- An agent embedded in the Cisco MDS 9000 Family switches. This agent can be selectively enabled to promiscuously capture designated control traffic.
- A text-based interface to the control and decoded output of the analyzer.
- GUI-based client application that you can install on any workstation to provide a full-function interface to the decoded data.

The text-based interface to the Cisco Fabric Analyzer is a CLI-based program for controlling the analyzer and providing output of the decoded results. Using the CLI-based interface, you can remotely access an Cisco MDS 9000 Family switch, using Telnet or a secure method such as Secure Shell (SSH). You can then capture and decode Fibre Channel control traffic, which offers a convenient method for conducting detailed, remote troubleshooting. In addition, because this tool is CLI-based, you can use roles-based policies to limit access to this tool as required.

The GUI-based implementation (Ethereal) can be installed on any Windows or Linux workstation. This application provides an easier-to-use interface that is more easily customizable. The GUI interface lets you easily sort, filter, crop, and save traces to your local workstation.

The Ethereal application allows remote access to Fibre Channel control traffic and does not require a Fibre Channel connection on the remote workstation.

The Cisco Fabric Analyzer lets you capture and decode Fibre Channel traffic remotely over Ethernet. It captures Fibre Channel traffic, encapsulates it in TCP/IP, and transports it over an Ethernet network to the remote client. The remote client then deencapsulates and fully decodes the Fibre Channel frames. This capability provides flexibility for troubleshooting problems in remote locations.

The Cisco Fabric Analyzer captures and analyzes control traffic coming to the Supervisor Card. This tool is much more effective than the debug facility for packet trace and traffic analysis, because it is not very CPU intensive and it provides a graphic interface for easy analysis and decoding of the captured traffic.

```
switch# config terminal
switch(config)# fcanalyzer local brief
Capturing on eth2
 0.000000 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x59b7 0xffff 0x7 -> 0xf HLO
 0.000089 ff.ff.fd -> ff.ff.fd FC 1 0x59b7 0x59c9 0xff -> 0x0 Link Ctl, ACK1
 1.991615 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x59ca 0xffff 0xff -> 0x0 HLO
 1.992024 ff.ff.fd -> ff.ff.fd FC 1 0x59ca 0x59b8 0x7 -> 0xf Link Ctl, ACK1

fcanalyzer example of fully decoded frame.
switch2(config)# fcanalyzer local
Capturing on eth2
Frame 1 (96 bytes on wire, 96 bytes captured)
  Arrival Time Jan 13, 2003 135038.787671000
  Time delta from previous packet 0.000000000 seconds
  Time relative to first packet 0.000000000 seconds
  Frame Number 1
  Packet Length 96 bytes
  Capture Length 96 bytes
Ethernet II, Src 00000000000a, Dst 00000000ee00
  Destination 00000000ee00 (00000000ee00)
```


Send comments to mdsfeedback-doc@cisco.com.

```

Source 00000000000a (00000000000a)
Type Vegas FC Frame Transport (0xfcfc)
MDS Header(SOFF/EOFn)
MDS Header
  Packet Len 66
  .... 0000 0001 11.. = Dst Index 0x0007
  .... ..00 1111 1111 = Src Index 0x00ff
  .... 0000 0000 0001 = VSAN 1
MDS Trailer
  EOF EOFn (3)
Fibre Channel
  R_CTL 0x02
  Dest Addr ff.fc.7e
  CS_CTL 0x00
  Src Addr ff.fc.7f
  Type SW_ILS (0x22)
  F_CTL 0x290000 (Exchange Originator, Seq Initiator, Exchg First, Seq Last,
CS_CTL, Transfer Seq Initiative, Last Data Frame - No Info, ABTS - Abort/MS, )
  SEQ_ID 0x11
  DF_CTL 0x00
  SEQ_CNT 0
  OX_ID 0x5a06
  RX_ID 0x0000
  Parameter 0x00000000
SW_ILS
  Cmd Code SW_RSCN (0x1b)
  0010 .... = Event Type Port is offline (2)
  .... 0000 = Address Format Port Addr Format (0)
  Affected Port ID 7f.00.01
  Detection Function Fabric Detected (0x00000001)
  Num Entries 1
  Device Entry 0
  Port State 0x20
  Port Id 7f.00.01
  Port WWN 1000000530005f1f (000530)
  Node WWN 1000000530005f1f (000530)

```

However, the Cisco Fabric Analyzer is not the right tool for troubleshooting end-to-end problems because it cannot access any traffic between the server and storage subsystems. That traffic is switched locally on the linecards, and does not reach the Supervisor card. In order to debug issues related to the communication between server and storage subsystems, you need to use Fibre Channel SPAN with an external protocol analyzer.

There are two ways you can start the Cisco Fabric Analyzer from the CLI.

- **fcanalyzer local**—Launches the text-based version on the analyzer directly on the console screen or on a file local to the system.
- **fcanalyzer remote *ip address***—Activates the remote capture agent on the switch, where *ip address* is the address of the management station running Ethereal.

CiscoWorks RME

CiscoWorks Resource Manager Essentials (RME) is a set of CiscoWorks applications that provide comprehensive resource management capabilities. With the introduction of the Cisco MDS 9000 Family of switches, CiscoWorks RME has been extended to provide resource management services to a Cisco MDS 9000 Family storage network.

CiscoWorks RME comprises a set of resource management services. The following list outlines the services provided by CiscoWorks RME for the Cisco MDS 9000 Family of switches.

Send comments to mdsfeedback-doc@cisco.com.

- **Inventory Manager**—Provides a facility to gather and audit a detailed hardware and software inventory of all Cisco MDS 9000 Family devices deployed in the storage network. A reporting facility is included to generate inventory reports
- **Configuration Manager**—Maintains an active repository of device configuration files for devices that are managed. It provides facility to upload and download configuration files to/from devices and a facility to log a record in the Change Audit log database when a new version of the configuration file is archived. Standard reports can be generated for configuration management inventory and activity.
- **Configuration Editor**—Provides a powerful web-based editor that allows multiple configuration files to be checked out of the configuration archive, be updated or changed, and then either saved locally or downloaded to the device.
- **Net Show**—Provides a simplified web-based show command interface, allowing show commands to be run against multiple switches or routers to enhance and simplify network troubleshooting.
- **Software Image Manager**—Simplifies version management and routine deployment of software updates to Cisco devices through wizard-assisted planning, scheduling, downloading, and monitoring of software updates
- **Syslog Analyzer**—Filters Syslog messages logged by Cisco devices and displays explanations of probable causes and recommended actions. This tool also helps facilitate manual parsing of Syslog files for reporting purposes.

CiscoWorks RME provides a system that can manage hardware, software, and configuration inventory across multiple infrastructures including storage networks, LANs, MANs, and WANs.

Using Other Troubleshooting Products

This section describes products from other vendors that you might find useful when troubleshooting problems with your storage network and connected devices. It includes the following topics:

- Fibre Channel Testers, page 1-20
- Fibre Channel Protocol Analyzers, page 1-20

Fibre Channel Testers

Fibre Channel testers are generally used to troubleshoot low-level protocol functions (such as Link Initialization). Usually these devices operate at 1- or 2-Gbps and provide the capability to create customized low-level Fibre Channel primitive sequences.

Fibre Channel testers are primarily used to ensure physical connectivity and low-level protocol compatibility, such as with different operative modes like Point-to-Point or Loop mode.

Fibre Channel testers and more generalized optical testers may be used to spot broken cables, speed mismatch, link initialization problems and transmission errors. These devices sometimes incorporate higher-level protocol analysis tools and may be bundled with generic protocol analyzers.

Fibre Channel Protocol Analyzers

An external protocol analyzer (for example from Finisar), is capable of capturing and decoding link level issues and the fibre channel ordered sets which comprise the fibre channel frame. The Cisco Port Analyzer Adapter, does not capture and decode at the ordered set level.

Send comments to mdsfeedback-doc@cisco.com.

A Fibre Channel protocol analyzer captures transmitted information from the physical layer of the Fibre Channel network. Because these devices are physically located on the network instead of at a software re-assembly layer like most Ethernet analyzers, Fibre Channel protocol analyzers can monitor data from the 8b/10b level all the way to the embedded upper-layer protocols.

Fibre Channel network devices (HBAs, switches, and storage subsystems) are not able to monitor many SAN behavior patterns. Also, management tools that gather data from these devices are not necessarily aware of problems occurring at the Fibre Channel physical, framing, or SCSI upper layers for a number of reasons.

Fibre Channel devices are specialized for handling and distributing incoming and outgoing data streams. When devices are under maximum loads, which is when problems often occur, the device resources available for error reporting are typically at a minimum and are frequently inadequate for accurate error tracking. Also, Fibre Channel host bus adapters (HBAs) do not provide the ability to capture raw network data.

For these reasons, a protocol analyzer may be more important in troubleshooting a storage network than in a typical Ethernet network. There are a number of common SAN problems that occur in deployed systems and test environments that are visible only with a Fibre Channel analyzer. These include the following:

- Credit starvation
- missing, malformed, or non-standard-compliant frames or primitives
- protocol errors

Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting Switch System Issues

This chapter describes how to identify and resolve problems that might occur when accessing or starting up a single Cisco MDS 9000 Family switch. It includes the following sections:

- Recovering the Administrator Password, page 2-1
- Troubleshooting System Restarts, page 2-1

Recovering the Administrator Password

If you forget the administrator password for accessing a Cisco MDS 9000 Family switch, you can recover the password using a local console connection. For the latest instructions on password recovery, go to <http://www.cisco.com/warp/public/474/> and click on “MDS 9000 Series Multilayer Directors and Fabric Switches” under Storage Networking Routers.

Troubleshooting System Restarts

This section describes the different types of system crashes and how to respond to each type. It includes the following topics:

- Overview, page 2-1
- Working with Unrecoverable System Restarts, page 2-5

Overview

There are three different types of system restarts:

- Recoverable—A process restarts and service is not affected.
- Unrecoverable—A process is not restartable or it has restarted more than the max restart times within a fixed period of time (seconds) and will not be restarted again.
- System Hung/Crashed—No communications of any kind is possible with box.

Most system restarts generate a Call Home event, but the condition causing a restart may become so severe that a Call Home event is not generated. Be sure that you configure the Call Home feature properly, follow up on any initial messages regarding system restarts, and fix the problem before it becomes so severe. For information about configuring Call Home, refer to the *Cisco MDS 9000 Family Configuration Guide* or the *Cisco MDS 9000 Family Fabric Manager User Guide*.

Send comments to mdsfeedback-doc@cisco.com.

Working with Recoverable Restarts

Every process restart generates a Syslog message and a Call Home event. Even if the event is not service affecting you should identify and resolve the condition immediately because future occurrences could cause service interruption.

To respond to a recoverable system restart, follow these steps:

- Step 1** Enter the following command to check the Syslog file to see which process restarted and why it restarted.

```
switch# sh log logfile | include error
```

For information about the meaning of each message, refer to the *Cisco MDS 9000 Family System Messages Guide*

The system output looks like the following:

```
Sep 10 23:31:31 dot-6 % LOG_SYSMGR-3-SERVICE_TERMINATED: Service "sensor" (PID 704) has
finished with error code SYSMGR_EXITCODE_SY.
switch# show logging logfile | include fail
Jan 27 04:08:42 88 %LOG_DAEMON-3-SYSTEM_MSG: bind() fd 4, family 2, port 123, ad
dr 0.0.0.0, in_classd=0 flags=1 fails: Address already in use
Jan 27 04:08:42 88 %LOG_DAEMON-3-SYSTEM_MSG: bind() fd 4, family 2, port 123, ad
dr 127.0.0.1, in_classd=0 flags=0 fails: Address already in use
Jan 27 04:08:42 88 %LOG_DAEMON-3-SYSTEM_MSG: bind() fd 4, family 2, port 123, ad
dr 127.1.1.1, in_classd=0 flags=1 fails: Address already in use
Jan 27 04:08:42 88 %LOG_DAEMON-3-SYSTEM_MSG: bind() fd 4, family 2, port 123, ad
dr 172.22.93.88, in_classd=0 flags=1 fails: Address already in use
Jan 27 23:18:59 88 % LOG_PORT-5-IF_DOWN: Interface fc1/13 is down (Link failure
or not-connected)
Jan 27 23:18:59 88 % LOG_PORT-5-IF_DOWN: Interface fc1/14 is down (Link failure
or not-connected)
Jan 28 00:55:12 88 % LOG_PORT-5-IF_DOWN: Interface fc1/1 is down (Link failure o
r not-connected)
Jan 28 00:58:06 88 % LOG_ZONE-2-ZS_MERGE_FAILED: Zone merge failure, Isolating p
ort fc1/1 (VSAN 100)
Jan 28 00:58:44 88 % LOG_ZONE-2-ZS_MERGE_FAILED: Zone merge failure, Isolating p
ort fc1/1 (VSAN 100)
Jan 28 03:26:38 88 % LOG_ZONE-2-ZS_MERGE_FAILED: Zone merge failure, Isolating p
ort fc1/1 (VSAN 100)
Jan 29 19:01:34 88 % LOG_PORT-5-IF_DOWN: Interface fc1/1 is down (Link failure o
r not-connected)
switch#
```

- Step 2** Enter the following command to identify the processes that are running and the status of each process.

```
switch# show processes
```

The following codes are used in the system output for the State (process state):

- D = uninterruptible sleep (usually IO)
- R = runnable (on run queue)
- S = sleeping
- T = traced or stopped
- Z = defunct ("zombie") process
- NR = not-running
- ER = should be running but currently not-running

Send comments to mdsfeedback-doc@cisco.com.

**Note**

ER usually is the state a process enters if it has been restarted too many times and has been detected as faulty by the system and disabled.

The system output looks like the following (the output has been abbreviated to be more concise):

PID	State	PC	Start_cnt	TTY	Process
-----	-----	-----	-----	----	-----
1	S	2ab8e33e	1	-	init
2	S	0	1	-	keventd
3	S	0	1	-	ksoftirqd_CPU0
4	S	0	1	-	kswapd
5	S	0	1	-	bdfldush
6	S	0	1	-	kupdated
71	S	0	1	-	kjournald
136	S	0	1	-	kjournald
140	S	0	1	-	kjournald
431	S	2abe333e	1	-	httpd
443	S	2abfd33e	1	-	xinetd
446	S	2ac1e33e	1	-	sysmgr
452	S	2abe91a2	1	-	httpd
453	S	2abe91a2	1	-	httpd
456	S	2ac73419	1	S0	vsh
469	S	2abe91a2	1	-	httpd
470	S	2abe91a2	1	-	httpd

- Step 3** Enter the following command to show the processes that have had abnormal exits and if there is a stack-trace or core dump.

```
switch# show process log
```

Process	PID	Normal-exit	Stack-trace	Core	Log-create-time
-----	-----	-----	-----	-----	-----
ntp	919	N	N	N	Jan 27 04:08
snsn	972	N	Y	N	Jan 24 20:50

- Step 4** Enter the following command to show detailed information about a specific process that has restarted:

```
switch# show processes log pid 898
```

The system output looks like the following:

```
Service: idehsd
Description: ide hotswap handler Daemon
Started at Mon Sep 16 14:56:04 2002 (390923 us)
Stopped at Thu Sep 19 14:18:42 2002 (639239 us)
Uptime: 2 days 23 hours 22 minutes 22 seconds
Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGTERM (3)
Exit code: signal 15 (no core)
CWD: /var/sysmgr/work
Virtual Memory:
CODE      08048000 - 0804D660
  DATA   0804E660 - 0804E824
  BRK     0804E9A0 - 08050000
  STACK   7FFFFD10

Register Set:
EBX 00000003      ECX 0804E994      EDX 00000008
  ESI 00000005      EDI 7FFFC9C      EBP 7FFFCAC
  EAX 00000008      XDS 0000002B      XES 0000002B
  EAX 00000003 (orig) EIP 2ABF5EF4      XCS 00000023
  EFL 00000246      ESP 7FFFC5C      XSS 0000002B

Stack: 128 bytes. ESP 7FFFC5C, TOP 7FFFFD10
0x7FFFC5C: 0804F990 0804C416 00000003 0804E994 .....
```

Send comments to mdsfeedback-doc@cisco.com.

```
0x7FFFC6C: 00000008 0804BF95 2AC451E0 2AAC24A4 .....Q.*.*
0x7FFFC7C: 7FFFD14 2AC2C581 0804E6BC 7FFFC8A8 .....*.....
0x7FFFC8C: 7FFFC94 00000003 00000001 00000003 .....
0x7FFFC9C: 00000001 00000000 00000068 00000000 .....h.....
0x7FFFCAC: 7FFFC8E 2AB4F819 00000001 7FFFD14 .....*.....
0x7FFFCBC: 7FFFD1C 0804C470 00000000 7FFFC8E ....p.....
0x7FFFC8C: 2AB4F7E9 2AAC1F00 00000001 08048A2C ...*...*.....
PID: 898
SAP: 0
UUID: 0
switch#
```

Step 5 Enter the following command to determine if the restart recently occurred:

```
switch# sh sys uptime
Start Time: Fri Sep 13 12:38:39 2002
Up Time: 0 days, 1 hours, 16 minutes, 22 seconds
```

To determine if the restart is repetitive or a one-time occurrence, compare the length of time that the system has been up with the timestamp of each restart.

Step 6 Enter the following command to view the core files:

```
switch# show cores
```

The system output looks like the following:

Module-num	Process-name	PID	Core-create-time
-----	-----	---	-----
5	fspf	1524	Jan 9 03:11
6	fcc	919	Jan 9 03:09
8	acltcam	285	Jan 9 03:09
8	fib	283	Jan 9 03:08

This output shows all the cores presently available for upload from the active supervisor. The column entitled *module-num* shows the slot# on which the core was generated. In the example shown above, an fspf core was generated on the active supervisor module in slot 5. An fcc core was generated on the standby supervisory module in slot 6. Core dumps generated on the line card in slot 8 include acltcam and fib.

To copy the FSPF core dump in this example to a TFTP server with the IP address 1.1.1.1, enter the following command:

```
switch# copy core://5/1524 tftp://1.1.1.1/abcd
```

The following command displays the file named `zone_server_log.889` in the `log` directory.

```
switch# sh pro log pid 1473
=====
Service: ips
Description: IPS Manager

Started at Tue Jan 8 17:07:42 1980 (757583 us)
Stopped at Thu Jan 10 06:16:45 1980 (83451 us)
Uptime: 1 days 13 hours 9 minutes 9 seconds

Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 6 (core dumped)
CWD: /var/sysmgr/work

Virtual Memory:

CODE      08048000 - 080FB060
DATA      080FC060 - 080FCBA8
```


Send comments to mdsfeedback-doc@cisco.com.

```
BRK      081795C0 - 081EC000
STACK    7FFFFCF0
TOTAL    20952 KB
```

Register Set:

```
EBX 000005C1      ECX 00000006      EDX 2AD721E0
ESI 2AD701A8      EDI 08109308      EBP 7FFFF2EC
EAX 00000000      XDS 0000002B      XES 0000002B
EAX 00000025 (orig) EIP 2AC8CC71      XCS 00000023
EFL 00000207      ESP 7FFFF2C0      XSS 0000002B
```

Stack: 2608 bytes. ESP 7FFFF2C0, TOP 7FFFFCF0

```
0x7FFFF2C0: 2AC8C944 000005C1 00000006 2AC735E2 D..*.....5.*
0x7FFFF2D0: 2AC8C92C 2AD721E0 2AAB76F0 00000000 ,...*!.*.v.*....
0x7FFFF2E0: 7FFFF320 2AC8C920 2AC513F8 7FFFF42C ... ..*...*,...
0x7FFFF2F0: 2AC8E0BB 00000006 7FFFF320 00000000 ...*.... ..
0x7FFFF300: 2AC8DF8 2AD721E0 08109308 2AC65AFC ...*!.*.....Z.*
0x7FFFF310: 00000393 2AC6A49C 2AC621CC 2AC513F8 .....*!.*...*
0x7FFFF320: 00000020 00000000 00000000 00000000 .....
0x7FFFF330: 00000000 00000000 00000000 00000000 .....
0x7FFFF340: 00000000 00000000 00000000 00000000 .....
0x7FFFF350: 00000000 00000000 00000000 00000000 .....
0x7FFFF360: 00000000 00000000 00000000 00000000 .....
0x7FFFF370: 00000000 00000000 00000000 00000000 .....
0x7FFFF380: 00000000 00000000 00000000 00000000 .....
0x7FFFF390: 00000000 00000000 00000000 00000000 .....
0x7FFFF3A0: 00000002 7FFFF3F4 2AAB752D 2AC5154C .
... output abbreviated ...
```

Stack: 128 bytes. ESP 7FFFF830, TOP 7FFFFCD0

Step 7 Enter the following command configure the switch to use TFTP to send the core dump to a TFTP server.

```
switch(config)# sys cores tftp://servername[/path]
```

This command causes the switch to enable the automatic copy of core files to a TFTP server. For example, the following command sends the core files to the TFTP server with the IP address 10.1.1.1.

```
switch(config)# system cores tftp://10.1.1.1/cores
```

The following conditions apply:

- The core files are copied every 4 minutes. This time is not configurable.
- The copy of a specific core file can be manually triggered, using the command **copy core//module#/pid# tftp//tftp_ip_address/file_name**
- The maximum number of times a process can be restarted is part of the HA policy for any process (this parameter is not configurable). If the process restarts more than the maximum number of times, the older core files are overwritten.
- The maximum number of core files that can be saved for any process is part of the HA policy for any process (this parameter is not configurable, and it is set to 3).

Step 8 To determine the cause and resolution for the restart condition, call Cisco TAC and ask them to review your core dump.

Working with Unrecoverable System Restarts

An unrecoverable system restart may occur in the following cases:

Send comments to mdsfeedback-doc@cisco.com.

- A critical process fails and is not restartable
- A process restarts more times than is allowed by the system configuration
- A process restarts more frequently than is allowed by the system configuration

The effect of a process restart is determined by the policy configured for each process. Unrecoverable restarts may cause loss of functionality, restart of the active supervisor, a supervisor switchover, or restart of the switch.

To respond to an unrecoverable restart, perform the steps listed in the “Working with Recoverable Restarts” section on page 2-2.

Troubleshooting Switch Level Issues and Interswitch Connectivity

This chapter describes how to identify and resolve problems that affect basic connectivity between switches, hosts, and storage in the network fabric.

The most common problems a system administrator can face may be categorized by two different scenarios:

- Switch-to-switch basic interconnectivity problems, which can result in the isolation of a port or VSAN due to incorrect parameters or settings on an ISL or VSAN
- Fabric to Server/Storage connectivity problems: identified by a Fx port not coming up or caused by zone or VSAN configuration errors

This section will present some of the most common scenarios, in which either switch-to-switch basic connectivity problems or fabric to end-devices problems can be found.

These scenario are grouped in three different sections:

- Troubleshooting E Port Connectivity - ISL Isolation, page 3-1
- Troubleshooting TE Port Connectivity - VSAN Isolation, page 3-10
- Troubleshooting Fx Port Connectivity, page 3-13

Troubleshooting E Port Connectivity - ISL Isolation

This section describes how to troubleshoot Inter-Switch Link (ISL) isolation, which may occur when you try to merge two separate fabrics or add a new domain to an existing fabric. It includes the following topics:

- Overview, page 3-2
- Troubleshoot Switch and Port Parameters, page 3-3
- Troubleshooting a ZONE Merge Failures, page 3-5
- Troubleshooting a VSAN Configuration Conflict, page 3-6
- Troubleshooting a Domain ID Conflict, page 3-7

Send comments to mdsfeedback-doc@cisco.com.

Overview

On an E port, only one VSAN can be passed and possibly be isolated. However, in a trunking E port (TE), multiple VSANs can become isolated while others are passing traffic. The same troubleshooting approach applies in both cases, except that on a trunking E port the troubleshooting may need to be done on a per-VSAN basis and/or on multiple VSANs.

Step 1) Verify that each VSAN is able to see the other switches within the same VSAN.

To verify that each switch is able to see the other switches, issue the following command at the exec prompt. This command is VSAN-specific. If a specific VSAN is omitted from the command, it will list the output for all VSANs.

```
switch# show fcdomain domain-list vsan 1
```

The output of the command lists set of domain IDs and associated WWNs for each switch within a VSAN. This list provides the WWN of the switches owning each domain ID and the information about the principal of those switches in the fabric or VSAN they belong to.

Sample output #1 (obtained in a fabric with just 2 switches in VSAN 1)

```
switch# sh fcdomain domain-list vsan 1

Number of domains: 2
Domain ID          WWN
-----
0x4a(74)           20:01:00:05:30:00:13:9f [Local]
0x4b(75)           20:01:00:05:30:00:13:9e [Principal]
-----
```

This is an output of VSAN 1 in which 2 switches are seen. This indicates that the switch where the command has been issued has built its adjacency on VSAN 1 with the other switch in the same VSAN.

Sample output #2

```
switch# sh fcdomain domain-list vsan 1

Number of domains: 1
Domain ID          WWN
-----
0x4a(74)           20:01:00:05:30:00:13:9f [Local] [Principal]
-----
```

In this output only one switch is seen. This indicates that the switch where the command has been issued has not established adjacency with the neighboring switch in VSAN 1.

Possible Causes for the problem

In the previous output one domain ID is missing. The reason of this can be found in the erroneous configuration of several parameters between adjacent switches.

In case of E port/VSAN isolation the following parameters and configurations should be checked in each switch:

- Zoning configuration
- VSAN configuration
- Domain ID parameters
- Fabric Parameters & Timers

Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting steps

The first step in the process is to understand the nature of the problem by checking the status of the E port. The **show interface** command can be used to display the port status. This command is very useful in troubleshooting switch connectivity issues. In case of error, this command provides indication of the protocol error or configuration mismatch that caused the problem, between parenthesis, immediately after the port operational status.

Sample output#3

```
switch# show interface fc4/1
fc4/1 is up
  Hardware is Fibre Channel
  Port WWN is 20:c1:00:05:30:00:13:9e
  Peer port WWN is 21:89:00:05:30:00:18:a2
  Admin port mode is auto, trunk mode is on
  Port mode is E, FCID is 0x6b0000
  Port vsan is 1
  Speed is 2 Gbps
  Receive B2B Credit is 12
  Receive Buffer Size is 2112
  Encapsulation is normal
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    109 frames input, 9728 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    108 frames output, 6652 bytes, 0 discards
    1 input OLS, 3 LRR, 1 NOS, 2 loop inits
    4 output OLS, 3 LRR, 3 NOS, 0 loop inits
```

When the port operational status shows that the port is up, it means that the E port is up and is currently passing traffic. The output shown above represents the output of a working E port.

Different types of isolation problems can be recognized by running the **show interface** command:

- Isolation due to ELP failure and mismatch in the switch or port parameters
- Isolation due to zone merge failure
- Isolation due to port/VSAN mismatch
- Isolation due to domain overlap
- Isolation due to invalid fabric reconfiguration

Troubleshoot Switch and Port Parameters

One of the possible causes for E port isolation is a mismatch in the configured switch or port parameters. The problem will affect the link initialization process, and eventually the initial ELP exchange. An example of the expected **show interface** command output is shown below:

```
switch# show interface fc2/4
fc2/4 is down (Isolation due to ELP failure)
  Hardware is Fibre Channel, WWN is 20:44:00:05:30:00:18:a2
  vsan is 1
  Beacon is turned off
  1445517676 packets input, 727667035658 bytes, 0 discards
  0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
```

Send comments to mdsfeedback-doc@cisco.com.

```
Received 0 runts, 0 jabber, 0 too long, 0 too short
    0 EOF abort, 0 fragmented, 0 unknown class
    100 OLS, 67 LRR, 37 NOS, 0 loop inits
```

```
133283352 packets output, 1332969530 bytes
Transmitted 198 OLS, 50 LRR, 0 NOS, 10 loop inits
```

In this example the interface is indicating a link isolation caused by an ELP failure on an E port. The ELP is a frame sent between two switches to negotiate fabric parameters. If you get this failure, verify the fabric parameters are the same for both switches.

Since fabric parameters are configured on a per switch basis, however, they are required to be the same for all switches within a fabric.

If the interface indicates an ELP failure verify the following parameters match using the **show fctimer** command:

- ED_TOV Timer
- RA_TOV Timer
- FS_TOV timer

An example of the show fctimer command where all default value are in use, is shown below

```
switch# show fctimer
F_S_TOV : 5000 milliseconds
D_S_TOV : 15000 milliseconds
E_D_TOV : 2000 milliseconds
R_A_TOV : 10000 milliseconds
```

Another parameter to check, is the Rcxbuffer size on the interface. This value should match on both the ends of a ISL. To verify the value of the Rcxbuffer size, use the following command:

```
switch# show port internal info interface fc2/1

fc2/1 - if_index: 1080000
Admin Config - state(up), mode(Auto), speed(auto), trunk(no trunk)
    beacon(off), snmp trap(on), tem(false)
    bb_credit(default), rxbufsize(2112), encap(default)
    description()
Operational Info - state(down), mode(ALL), speed(auto), trunk(no trunk)
    state reason(Link failure or not-connected)
    phy port enable (1), phy layer (FC)
    participating(1), port_vsan(1), null_vsan(0), fcid(0x000000)
    current state [PI_FSM_ST_LINK_INIT]
    port_init_eval_flag(0x00000001), cfg wait for none
    Mts node id 0x202
    cnt_link_failure(0), cnt_link_success(0), cnt_port_up(0)
    cnt_cfg_wait_timeout(0), cnt_port_cfg_failure(0), cnt_init_retry(0)
Port Capabilities -
    Modes: E,TE,F,FL,TL,SD
    Min Speed: 1000
    Max Speed: 2000
    Max Tx Bytes: 2112
    Max Rx Bytes: 2112
    Max Tx Buffer Credit: 255
    Max Rx Buffer Credit: 16
    Max Private Devices: 63
    Max Sourcedable Pkt Size: 2112
    Hw Capabilities: 0xb
    Connector Type: 0x0
```

Send comments to mdsfeedback-doc@cisco.com.

```

FCOT info -
  Min Speed: 1000
  Max Speed: 2000
  Module Type: 8
  Connector Type: 7
  Gigabit Eth Compliance Codes: 0
  FC Transmitter Type: 3
  Vendor Name: PICOLIGHT
  Vendor ID: 0:4:133
  Vendor Part Num: PL-XPL-00-S23-28
  Vendor Revision Level:
Trunk Info -
  trunk vsans (allowed active) (1)

```

In the above examples, the highlighted rxbuffsize is 2112 bytes. This represents the default settings on a Cisco MDS 9000 switch interface.

Troubleshooting a ZONE Merge Failures

In the example below the **show interface** command indicates that the E port did not come up due to a zone merge failure. (Zoning information is on a per VSAN basis. Therefore, for a TE port, it may be necessary to verify that the zoning information does not conflict for any allowed VSAN.)

```

switch# show interface fc2/14

fc2/14 is down (Isolation due to zone merge failure)
  Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e
  vsan is 1
  Beacon is turned off
    40 frames input, 1056 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 3 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    79 frames output, 1234 bytes, 16777216 discards
    Received 23 OLS, 14 LRR, 13 NOS, 39 loop inits
    Transmitted 50 OLS, 16 LRR, 21 NOS, 25 loop inits

```

A E port will segment with (isolation due to zone merge failure) if the following are true.

- If the active zoneset is different on the two switches.
- If the active zoneset on both switches contain a zone with the same name but with different zone members

To verify the zoning information use the following commands:

```

switch# show zone vsan 1
switch# show zoneset vsan 1

```

Two different approaches may be followed to solve a zone merge failure:

- Overwrite the zoning configuration of one switch with the other switch's configuration. This can be done with the following commands:

```

switch# zone merge interface fc2/7 import vsan 1
switch# zone merge interface fc2/7 export vsan 1

```

The import option of the command will overwrite the local switch's active zoneset with that of the remote switch. The export command will overwrite the remote switch's active zoneset with the local switch's active zoneset.

Send comments to mdsfeedback-doc@cisco.com.

- If the zoning databases between the two switches are overwritten, this option is not possible (because the final configuration should be a combination of the configuration of the two switches). Another option is to manually change the content of the zone database on either of the switches, and then have another cycle up/down issued on the isolated port.

If the isolation is specific to one VSAN and not on a E port, the correct way to issue the cycle up/down, is to remove the VSAN from the list of allowed VSANs on that trunk port, and re-insert it.

Do not simply issue a **shut/no shut** sequence of commands on the port, because this would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the isolation problem.

Using the Zone Merge Analysis tool in Fabric Manager, the compatibility of two active zonesets in two switches can be checked before actually merging the two zonesets. Refer to the *Cisco MDS 9000 Fabric Manager User Guide* for more information.

Troubleshooting a VSAN Configuration Conflict

In the following example, the E port has been isolated because the interfaces connecting the two switches belong to different VSANs.

```
switch# show interface fc2/4
fc2/4 is down fc2/4 is down (isolation due to port vsan mismatch)
```

```
Hardware is Fibre Channel, WWN is 20:44:00:05:30:00:63:5e
vsan is 2
Beacon is turned off
 30 frames input, 682 bytes, 0 discards
 0 runts, 0 jabber, 0 too long, 0 too short
 0 input errors, 0 CRC, 0 invalid transmission words
 0 address id, 0 delimiter
 0 EOF abort, 0 fragmented, 0 unknown class
 30 frames output, 583 bytes, 0 discards
Received 2 OLS, 2 LRR, 2 NOS, 5 loop inits
Transmitted 5 OLS, 3 LRR, 2 NOS, 4 loop inits
```

In the above example, the E port isolated because the interface of the 2 switches belong to different VSANS. To resolve this issue move the interfaces to the same VSAN.

You can check the VSAN membership with the switch interfaces with the following command.

```
switch# show vsan membership
vsan 1 interfaces:
    fc2/1  fc2/2  fc2/3  fc2/4  fc2/6  fc2/7  fc2/8  fc2/9
    fc2/10 fc2/11 fc2/12 fc2/14 fc2/15 fc2/16 fc7/1  fc7/2
    fc7/3  fc7/4  fc7/5  fc7/6  fc7/7  fc7/8  fc7/9  fc7/10
    fc7/11 fc7/12 fc7/13 fc7/14 fc7/15 fc7/16 fc7/17 fc7/18
    fc7/19 fc7/20 fc7/21 fc7/22 fc7/23 fc7/24 fc7/25 fc7/26
    fc7/27 fc7/28 fc7/29 fc7/30 fc7/31 fc7/32

vsan 2 interfaces:
    fc2/5  fc2/13

vsan 4094(isolated_vsan) interfaces:
```

The command shows that all the interfaces on the switch belong to VSAN 1, with the exception of interface 2/5 and 2/13 that are part of VSAN 2.

Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting a Domain ID Conflict

In a Fibre Channel network, the principal switch is used to issue domain IDs when a new switch is added to an existing fabric. However, when two fabrics merge, the principal switch selection process determines which one of the pre-existing switches becomes that principal switch. The election of the new principal switch is characterized by the following rules:

1. A switch with a non-empty domain ID list has priority over a switch that has an empty domain ID list, and the principal switch will be the principal switch of the first fabric. In the case of a single switch fabric, it does not contain a domain ID list.
2. If both fabrics have a domain ID list, the priority between the two principal switches is determined by configured switch priority. This is a user-settable parameter - the lower the value the higher the priority.
3. If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the World Wide Names of the two switches. The lower value has the higher priority.

When merging two fabrics, the administrator can expect the following behavior:

- When connecting a single-switch fabric to a multi-switch fabric, the multi-switch fabric always retains its principal switch regardless of the principal switch priority setting on the single switch fabric.
- When powering up a new switch that is connected to an existing fabric with two or more switches, the existing switch fabric always retains its principal switch, even if the new switch has a higher administratively assigned principal switch priority.
- When powering up a new switch that is connected to a standalone switch, the new principal switch is determined by the administratively set priority. If no priority is set (the default priority is used in every switch), it is determined by the World Wide Name. This also applies to connecting to two single-switch fabrics.
- When connecting a multi-switch fabric to another multi-switch fabric, the principal switch is determined by the administratively set priority. If no priority is set (default value is used by every switch) it is determined by the World Wide Name of the existing principal switches of the two fabrics.

There are several reasons why two switch fabrics would not merge:

- If two switch fabrics with two or more switches are connected, and both fabrics that have switches with the domain ID already assigned and the auto-reconfigure option is disabled (this option is disabled by default), the E ports that are used to connect the two fabrics will be isolated.

In this case, the following error message is returned in the **show interface** command.

```
switch# show interface fc2/14
fc2/14 is down (Isolation due to domain overlap)
  Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e
  vsan is 2
  Beacon is turned off
    192 frames input, 3986 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 3 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    231 frames output, 3709 bytes, 16777216 discards
    Received 28 OLS, 19 LRR, 16 NOS, 48 loop inits
    Transmitted 62 OLS, 22 LRR, 25 NOS, 30 loop inits
```

To view which domain are currently in your fabric use the following command. This command can be issued to both fabrics to determine which domain IDs are overlapping.

Send comments to mdsfeedback-doc@cisco.com.

```
switch# sh fcdomain domain-list vsan 1

Number of domains: 2
Domain ID          WWN
-----
0x4a(74)          20:01:00:05:30:00:13:9f [Local]
0x4b(75)          20:01:00:05:30:00:13:9e [Principal]
-----
```

The output shown is representative of a simple two switches fabric.

If a domain is currently isolated due to domain overlap, and you later enable the auto-reconfigure option on both switches, the fabric continues to be isolated. However, if you enable the auto-reconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. The RCF functionality would automatically force a new principal switch selection and cause new domain ID to be assigned to the different switches. A disruptive reconfiguration may affect data traffic.

To enable the auto-reconfigure option on a particular VSAN use the following command:

```
switch(config)# fcdomain auto-reconfigure vsan 10
```

There are different ways of resolving this issue. One possible solution is to cause an RCF in the fabric with a disruptive restart of the domain manager on the switch. The RCF functionality will cause all switches in both fabrics to empty their domain ID list and select a new principal switch. The overlapping domain ID would be assigned to which ever switch requests it first from the newly selected principal switch.

The second switch to request that same domain ID would be assigned a new domain ID. This will only work if the overlapping domain IDs are not statically defined on the switches. If the domain ID is statically defined on the switch, then the switch does not accept any other domain ID than the one that is configured. If its request for the configured domain ID is not granted, it isolates itself from the fabric. The RCF functionality is disruptive and can cause end nodes to be assigned new domain IDs. By default, Cisco MDS 9000 Family switches accept disruptive restarts. You can configure the switch in order to reject incoming RCFs on a per-VSAN and port level basis (enabling the **rcf-reject** option).

In case the **rcf-reject** option is on, in order to have the RCF propagated to the switches in the fabric, must have the RCF reject property disabled. To turn this option off use the following command:

```
switch(config-if)# no fcdomain rcf-reject vsan 1
```

At this point a disruptive restart should be triggered, by using the following command:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# fcdomain restart disruptive vsan 1
switch(config)#
```

When you do a disruptive restart on your switch and the switch on the other side does not have disruptive restart enabled the error on the E port would be the following

```
switch# show interface fc2/5
fc2/5 is down (isolation due to invalid fabric reconfiguration)
  Hardware is Fibre Channel
  Port WWN is 20:45:00:05:30:00:18:a2
  Admin port mode is auto, trunk mode is on
  Port vsan is 1
  Receive data field size is 2112
```

Send comments to mdsfeedback-doc@cisco.com.

```
Beacon is turned off
5 minutes input rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
342 frames input, 9492 bytes, 199 discards
0 CRC, 0 unknown class
0 too long, 0 too short
143 frames output, 4184 bytes, 0 discards
2 input OLS, 4 LRR, 2 NOS, 12 loop inits
6 output OLS, 2 LRR, 5 NOS, 1 loop inits
switch#
```

If the overlapping domain IDs are statically assigned, or you want to manually assign the new domain IDs, it would be necessary to go into the switch with the overlapping domain ID, configure a new domain ID for that switch, and restart the domain manager process to merge the two fabrics. To configure a domain ID on the switch use the following command. Domain IDs are assigned on a per-VSAN basis.

```
switch(config)# fcdomain domain 3 preferred vsan 1
switch(config)# fcdomain domain 3 static vsan 1
switch(config)#
```

When configuring a domain ID on a switch, there are two options - static and preferred. The static option tells the switch to request that particular domain in the fabric. If it does get that particular address, it will isolate itself from the fabric. With the preferred option, the switch requests the domain ID, but if that ID is not available it will accept another ID. After configuring the domain ID it is necessary to restart the domain manager to merge the two fabrics.

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# fcdomain restart vsan 1
switch(config)#
```

If a switch cannot get a statically configured address it would isolate itself from the fabric. When it isolates itself the **show interface** output is the following.

```
switch# show interface fc2/14
fc2/14 is down (isolation due to domain overlap)
Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e
vsan is 2
Beacon is turned off
192 frames input, 3986 bytes, 0 discards
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 3 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
231 frames output, 3709 bytes, 16777216 discards
Received 28 OLS, 19 LRR, 16 NOS, 48 loop inits
Transmitted 62 OLS, 22 LRR, 25 NOS, 30 loop inits
```

While of the neighbor switch E port the message would be the following

```
switch2# show interface fc9/4
fc9/4 is down (Isolation due to domain other side eport isolated)
Hardware is Fibre Channel
Port WWN is 22:04:00:05:30:00:13:9e
Admin port mode is auto, trunk mode is off
Port vsan is 1
Receive data field size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
```

Send comments to mdsfeedback-doc@cisco.com.

```

2547 frames input, 173332 bytes, 0 discards
  0 CRC,  0 unknown class
  0 too long, 0 too short
2545 frames output, 120300 bytes, 0 discards
14 input OLS, 12 LRR, 9 NOS, 33 loop inits
31 output OLS, 17 LRR, 13 NOS, 11 loop inits
switch2#

```

To fix the issue, the administrator can either change the static domain ID that is overlapping by manually configuring a new static domain ID for the isolated switch, or disable the static domain assignment, and allow the switch to request a new domain ID after a fabric reconfiguration.

For example, two switches may be merging and one of the switches is already powered up, but the second switch has a domain ID preconfigured. In this case, the switch that is powered up will have the priority in keeping the conflicting domain ID. The presently powered-on switch, being the principal switch, will assign a new domain ID to the new switch.

The E port could still isolate if the new switch has the domain ID configured as static, thus it would not take any other domain ID except for the one configured. If this were the case, it would be necessary to configure a new domain ID on the switch before the two switches could merge.

Troubleshooting TE Port Connectivity - VSAN Isolation

Trunking E ports (TE ports) are similar to E ports except they carry traffic for multiple VSANs. E ports carry traffic for a single VSAN. Because TE ports carry traffic for multiple VSANs, ISL isolation can affect one or more VSANs. For this reason, on a TE port you must troubleshoot ISL isolation on each VSAN.

Even in case of VSAN isolation, the starting point for troubleshooting the problem is to issue the **show interface** command.

The following example shows what the result of the **show interface** command, issued on a TE port in case of no VSAN is isolated (normal behavior).

```

switch# show interface fc7/9
fc7/9 is trunking
  Hardware is Fibre Channel
  Port WWN is 21:89:00:05:30:00:18:a2
  Peer port WWN is 20:c1:00:05:30:00:13:9e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  Speed is 2 Gbps
  vsan is 1
  Beacon is turned off
  FCID is 0x000000
  Receive B2B Credit is 12
  Trunk vsans (allowed active) (1,333)
  Trunk vsans (operational)    (1,333)
  Trunk vsans (up)            (1,333)
  Trunk vsans (isolated)      ()
  Trunk vsans (initializing)  ()
  Counter Values (current):
    262 frames input, 18808 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 2 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    264 frames output, 16960 bytes, 0 discards
    Received 5 OLS, 3 LRR, 2 NOS, 11 loop inits
    Transmitted 8 OLS, 4 LRR, 2 NOS, 3 loop inits

```

Send comments to mdsfeedback-doc@cisco.com.

```

Counter Values (5 minute averages):
  0 frames input, 136 bytes, 0 discards
  0 runts, 0 jabber, 0 too long, 0 too short
  0 input errors, 0 CRC, 0 invalid transmission words
  0 address id, 0 delimiter
  0 EOF abort, 0 fragmented, 0 unknown class
  0 frames output, 117 bytes, 0 discards
  Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
  Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
switch#

```

In the above example, the TE port is carrying traffic for VSANs 1 and 333. This is an example of a working TE port. There are no isolated VSANs in the list.

The following example shows the output of the **show interface** command, if one or more VSANs are isolated:

```

switch# show interface fc2/14
fc2/14 is trunking
  Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e
  Port mode is TE
  Speed is 2 Gbps
  vsan is 2
  Beacon is turned off
  Trunk vsans (allowed active) (1-3,5)
  Trunk vsans (operational)    (1-3,5)
  Trunk vsans (up)            (2-3,5)
  Trunk vsans (isolated)      (1)
  Trunk vsans (initializing)  ()
    475 frames input, 8982 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 3 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    514 frames output, 7509 bytes, 16777216 discards
    Received 30 OLS, 21 LRR, 18 NOS, 53 loop inits
    Transmitted 68 OLS, 25 LRR, 28 NOS, 32 loop inits

```

In the above example, the TE port has one VSAN isolated. The reason for VSAN isolation can be checked using the following command:

```
switch# sh interface fc2/14 trunk vsan 1
```

The command provides the same messages given by the **show interface**, in case of E port isolation. For example:

```

switch# sh interface fc2/14 trunk vsan 1
fc2/15 is trunking
  Vsan 1 is down (Isolation due to zone merge failure)

```

This output shows that VSAN 1 is isolated because of a zone merge error.

An alternative way to determine the cause of VSAN isolation, is to issue the following command:

```
switch#show port internal info interface fc2/14
```

The last few lines provide a description of the reason for VSAN isolation, for all the isolated VSANs.

```

switch# show port internal info int fc2/14

fc2/14 - if_index: 0x0109C000, phy_port_index: 0x3c
  Admin Config - state(up), mode(TE), speed(auto), trunk(on)

```

Send comments to mdsfeedback-doc@cisco.com.

```

beacon(off), snmp trap(on), tem(false)
rx bb_credit(default), rx bb_credit multiplier(default)
rxbufsize(2112), encap(default), user_cfg_flag(0x3)
description()
Operational Info - state(trunking), mode(TE), speed(2 Gbps), trunk(on)
state reason(None)
phy port enable (1), phy layer (FC)
participating(1), port_vsan(7), fcid(0x000000)
rx bb_credit multiplier(0), rx bb_credit(12)
current state [PI_FSM_ST_TEPORIT_INIT_TRUNKING_ENABLED]
port_init_eval_flag(0x00000001), cfg wait for none
Mts node id 0x202
eport_init_flag(0x00000386), elp_chk_flag(0x0000000A)
elp_rcvd_fc2_handle(0x00000000), elp_sent_fc2_handle(0x0800864E)
esc_chk_flag(0x0000002A), esc_fc2_handle(0x0801864F)
elp_flags(0x0000), classes_supported(F,2,3), tx bb_credit(255)
cnt_link_failure(1), cnt_link_success(5), cnt_port_up(0)
cnt_cfg_wait_timeout(0), cnt_port_cfg_failure(0), cnt_init_retry(0)
oper trunk mode(on)
Port Capabilities -
Modes: E,TE,F,FL,TL,SD
Min Speed: 1000
Max Speed: 2000
Max Sourcable Pkt Size: 2112
Max Tx Bytes: 2112
Max Rx Bytes: 2112
Max Tx Buffer Credit: 255
Max Rx Buffer Credit: 12
Max Rx Buffer Credit (ISL): 12
Default Rx Buffer Credit: 12
Default Rx Buffer Credit(ISL): 12
Default Rx Buffer Credit Multiplier: 0
Rx Buffer Credit change not allowed
Max Private Devices: 63
Hw Capabilities: 0xb
Connector Type: 0x0
FCOT info -
Min Speed: 1000
Max Speed: 2000
Module Type: 8
Connector Type: 7
Gigabit Eth Compliance Codes: 0
FC Transmitter Type: 3
Vendor Name: IBM
Vendor ID: 8:0:90
Vendor Part Num: IBM42P21SNY
Vendor Revision Level: AA20
Trunk Info -
trunk vsans (allowed active) (1,7-8)
trunk vsans (up) (7)
trunk vsans (isolated) (1,8)
TE port per vsan information
fc2/29, Vsan 1 - state(down), state reason(Isolation due to domain other side eport
isolated), fcid(0x000000)
port init flag(0x10000), current state [TE_FSM_ST_ISOLATED_DM_ZS]
fc2/29, Vsan 7 - state(up), state reason(None), fcid(0x690202)
port init flag(0x38000), current state [TE_FSM_ST_E_PORT_UP]
fc2/29, Vsan 8 - state(down), state reason(Isolation due to vsan not configure
d on peer), fcid(0x000000)
port init flag(0x0), current state [TE_FSM_ST_ISOLATED_VSAN_MISMATCH]

switch#

```

Send comments to mdsfeedback-doc@cisco.com.

In the example, VSAN 7 is up, while two VSANs are isolated (VSAN 1 and 8, the first one because of domain ID misconfiguration, and the second one because of VSAN misconfiguration).

Troubleshooting Fx Port Connectivity

This section describes how to troubleshoot problems with a switch interface in fabric mode (F port). An F port may be connected to a single N port, which is the mode used by peripheral devices (hosts or storage).

Two different major scenarios can be recognized in all the possible cases an administrator can incur in troubleshooting an Fx port:

- The port doesn't come up (just check the configuration of the interface, the cabling and the port connected to the switch).
- The port comes up, but the host is not able to communicate with the storage subsystem (in this case the configuration of VSAN and zones need to be checked)

Fx Port Fails to Achieve Up State

This section describes the troubleshooting steps to follow if, after connecting an N port to the switch, the port does not go into the up state.

In order for an Fx port to come up on a switch, the switch port must first acquire bit and word synchronization with the N port, and receive the FLOGI issued by the connected N port.

If any one of these steps does not occur, the switch port will not come up as an F port.

One of the first steps in troubleshooting an F port not coming up is to issue the **show interface** command from the CLI. This tells you where the process of interface initialization stops, and which steps to be used to solve the problem.

The following examples show different states of the interface, as displayed by the **show interface** command. In the examples, note the description of the problem that is printed immediately after the operational state of the interface.

FCOT Is Not Present

In the example below, the state of the interface is (Fcot is not present). This indicates that the switch does not detect the presence of an SFP on the interface. If this is the case verify that the SFP on the interface is seated properly. If reseating the SFP does not resolve the issue, replace the SFP or try another port on the switch.

```
switch# show interface fc7/31
fc7/31 is down (Fcot not present)
  Hardware is Fibre Channel
  Port WWN is 21:9f:00:05:30:00:18:a2
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Beacon is turned off
  Counter Values (current):
    0 frames input, 0 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
```

Send comments to mdsfeedback-doc@cisco.com.

```

0 frames output, 0 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Counter Values (5 minute averages):
0 frames input, 0 bytes, 0 discards
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 0 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
0 frames output, 0 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

```

Link_Failure or Not Connected

The following example shows the interface down as a result of a link failure or disconnect.

```

switch# show interface fc7/31
fc7/31 is down (Link_Failure or not connected)
Hardware is Fibre Channel
Port WWN is 21:9f:00:05:30:00:18:a2
Admin port mode is auto, trunk mode is on
vsan is 1
Beacon is turned off
Counter Values (current):
0 frames input, 0 bytes, 0 discards
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 0 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
0 frames output, 0 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Counter Values (5 minute averages):
0 frames input, 0 bytes, 0 discards
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 0 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
0 frames output, 0 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

```

The message “Link_Failure” or “not connected” can happen if nothing is connected to the specific interface (as in the case of a broken fiber), or if there is no bit synchronization between the switch interface and the Nx port directly connected. (In the case of a broken fiber, if **only** the TX path from the F port to the N port is broken, the switch interface will still have an operational Rx path, and so will still obtain bit synchronization from the bit stream received from the N port. It will also be able to recognize an incoming NOS from the N port and reply with an OLS. But because the transmitted OLS never reaches the N port the R_T_TOV timer expires. In this scenario the status of the port will also show “Link_Failure or not connected”.)

The key difference between this case and the “no bit synchronization” case, is that the input and output counts for OLS & NOS increment (as there is bit synchronization but no word synchronization). In such a state, you can check that the Tx path from switch to the Rx input on N port interface is properly connected. A faulty transmitter on the switch’s SFP or a faulty receiver on the N port’s SFP could also cause the issue.

Send comments to mdsfeedback-doc@cisco.com.

If, on the other hand, just the Rx path to the switch was broken and the TX path was intact, there would be no bit synchronization. In this case, the switch would not attempt to move to the link initialization stage, and therefore the OLS & NOS counters would never increment. Some possible causes of the problem could be a faulty SFP, or a speed mismatch between the Fx port and the Nx port, in case speed autonegotiation is disabled or doesn't work properly with the specific HBA/Storage port.

One of the first steps in the effort to solve the incompatibility, is to verify the speed configuration on the switch port. If the switch port is configured for a specific speed, configure the switch port for auto-speed detection. If the port still does not come up or the speed autonegotiation was already in place, the next possible step is to statically configure the speed of the port in order to match the one of the Nx port directly connected.

In case this doesn't solve the issue, and the port stays in the same status, the problem is probably a physical issue. Verify all physical parts between and including the switch port, the SFP on the switch port, and if used on the HBA, the HBA itself and the fiber connections.

Interface Bouncing between Offline and Initializing

In the above example, the state of the link bounces between initializing and offline. This indicates that the link has acquired bit and word synchronization, but has not been able to negotiate the type of link the switch port needs to become operative (because the initialization failed or an FLOGI is not issued by the HBAs or the FLOGI has been rejected by the switch).

```
switch# show interface fc7/2
fc7/2 is down (Initializing)
  Hardware is Fibre Channel
  Port WWN is 21:82:00:05:30:00:18:a2
  Admin port mode is F, trunk mode is on
  vsan is 1
  Beacon is turned off
  Counter Values (current):
    143274267 frames input, 182897329172 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    3016111132 frames output, 6029792843852 bytes, 0 discards
    Received 14 OLS, 14 LRR, 0 NOS, 31 loop inits
    Transmitted 87 OLS, 23 LRR, 65 NOS, 37 loop inits
  Counter Values (5 minute averages):
    0 frames input, 41 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    0 frames output, 112 bytes, 0 discards
    Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
    Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

switch# sh int fc7/2
fc7/2 is down (Offline)
  Hardware is Fibre Channel
  Port WWN is 21:82:00:05:30:00:18:a2
  Admin port mode is F, trunk mode is on
  vsan is 1
  Beacon is turned off
  Counter Values (current):
    143274267 frames input, 182897329172 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
```

Send comments to mdsfeedback-doc@cisco.com.

```

0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
3016111132 frames output, 6029792843852 bytes, 0 discards
Received 14 OLS, 14 LRR, 0 NOS, 31 loop inits
Transmitted 87 OLS, 23 LRR, 65 NOS, 37 loop inits
Counter Values (5 minute averages):
0 frames input, 41 bytes, 0 discards
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 0 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
0 frames output, 112 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

```

Usually a port in this state cycles between the initializing state and the off-line state. If this the error, verify that the switch port is configured as an F port. A port that is mistakenly set to be E port could cause this error when connecting to an N port. If setting the switch port to autodetect mode still doesn't resolve the issue, try to statically configure the port type in order to match the Nx port behavior with respect to loop or point-to-point initialization. There are some limited case in which the port type auto negotiation could not work properly for incompatibility with some old HBAs. (The automatic detection for the port mode is not applicable on an OSM module where at maximum one TE/E port can be configured per quad.)

If after having statically configured the port mode, it still doesn't come up, verify that the HBA is working properly and possibly power-cycle the HBA (i.e. booting the server or resetting the adapter).

To check whether a FLOGI is issued to the switch during the Fx port initialization, the **fc analyzer** can be used.

In the example below, the fc analyzer tool has been used with the necessary filtering to capture the FLOGI sent by the N port. Refer to the *Cisco MDS 9000 Family Configuration Guide* for more information.

```

Frame 25 (176 bytes on wire, 176 bytes captured)
Ethernet II
MDS Header(SOFi3/EOft)
Fibre Channel
FC ELS
  Cmd Code: FLOGI (0x04)
  Common Svc Parameters
    B2B Credit: 3
    Common Svc Parameters: 0x0 (Normal B2B Credit Mgmt)
    0000 .... = BB_SC Number: 0
    .... 1000 0100 0000 = Receive Size: 2112
    Max Concurrent Seq: 0
    Relative Offset By Info Cat: 0
    E_D_TOV: 0
    N_Port Port_Name: 10:00:00:09:c9:28:c7:01 (00:09:c9)
    Fabric/Node Name: 10:00:00:09:c9:28:c7:01 (00:09:c9)
  Class 1 Svc Parameters
    Service Options: 0x0 (Class Not Supported)
  Class 2 Svc Parameters
    Service Options: 0x0 (Class Not Supported)
  Class 3 Svc Parameters
    Service Options: 0x8800 (Seq Delivery Requested)
    Initiator Control: 0x0 (Seq Delivery Requested)
    Recipient Control: 0x0 (Seq Delivery Requested)
    Class Recv Size: 0
    Total Concurrent Seq: 0
    End2End Credit: 0
    Open Seq Per Exchg: 0
  Class 4 Svc Parameters

```

Send comments to mdsfeedback-doc@cisco.com.

```
Service Options: 0x0(Class Not Supported)
Vendor Version: 00000000000000000000000000000000
```

If the FLOGI is received by the switch, but the port is still not coming up, further investigation is needed to determine if the FLOGI is rejected by the switch (i.e. still using the fcanalyzer) or the cause is a misbehaving host, a broken fiber or a flappy connection between the end device and the Fx port on the switch. (If an Nx port is considered faulty by the driver/firmware or the ASIC used on the HBA, it can be configured to be in optical bypass. This results in the RX and TX paths being internally connected in the loopback by the on-board circuitry. If this happens, the switch port connected to the faulty device will reach bit and word synchronization with itself. If the port is configured in auto mode, this will cause the port to issue an ELP and try to initialize as a TE/E Port, even if an end device is physically connected to that interface. In this case a port reason code of isolation due to ELP failure, can show up even if an ISL is not present. To fix the issue, one possible approach is to reset the HBA or changing it if the problem persists.)

Point-to-point link comes up as FL_Port

If a point-to-point link comes up as a FL port, it could be caused by any of the following issues.

The switch port could be configured as either an autodetect port, or forced to be F port. In order to verify the actual configuration of the port, again the **show interface** command can be used:

```
switch#show interface fc7/5
fc7/5 is up
  Hardware is Fibre Channel
  Port WWN is 20:4d:00:05:30:00:18:a2
  Admin port mode is auto, trunk mode is on
  Port mode is F, FCID is 0x660000
  Port vsan is 1
  Speed is 1 Gbps
  Receive B2B Credit is 16
  Receive data field size is 2112
  Beacon is turned off
  5 minutes input rate 256 bits/sec, 32 bytes/sec, 0 frames/sec
  5 minutes output rate 256 bits/sec, 32 bytes/sec, 0 frames/sec
  369288 frames input, 11823952 bytes, 0 discards
    0 CRC, 0 unknown class
    0 too long, 0 too short
  369288 frames output, 11826732 bytes, 0 discards
  1 input OLS, 0 LRR, 2 NOS, 0 loop inits
  3 output OLS, 3 LRR, 1 NOS, 0 loop inits
```

An alternative way to get the same information and also the information about the configured speed is to use the following command.

```
switch# show port internal info interface fc7/5
```

The second line of the long output generated by the command shows the administrative status of the port (that reflect what configured by the administrator on the switch).

For example, issuing the command on a switch interface fc7/5 where a no shutdown command has been issued and where the port mode has been configured to be auto-detected, with speed auto-negotiation enabled and trunking capabilities enabled, the following output is displayed:

```
switch# show port internal info interface fc7/5
fc7/5 - if_index: 0x 1304000, phy_port_index: 0x84
  Admin Config - state(up), mode(auto), speed(auto), trunk(on)
  beacon(off), snmp trap(on), tem(false)
  rx_bb_credit(default), rx_bb_credit multiplier(default)
  rxbufsize(2112), encap(default), user_cfg_flag(0x1)
```

Send comments to mdsfeedback-doc@cisco.com.

```

description()
Operational Info - state(up), mode(F), speed(1 Gbps), trunk(off)
state reason(None)
phy port enable (1), phy layer (FC)
participating(1), port_vsan(1), null_vsan(0), fcid(0xef0300)
rx bb_credit multiplier(0), rx bb_credit(12)
current state [PI_FSM_ST_F_PORT_UP]
port_init_eval_flag(0x00003001), cfg wait for none
Mts node id 0x702
cnt_link_failure(54), cnt_link_success(53), cnt_port_up(4)
cnt_cfg_wait_timeout(0), cnt_port_cfg_failure(0), cnt_init_retry(0)
Port Capabilities -
Modes: E,TE,F,FL,TL,SD
Min Speed: 1000
Max Speed: 2000
Max Sourcingable Pkt Size: 2112
Max Tx Bytes: 2112
Max Rx Bytes: 2112
Max Tx Buffer Credit: 255
Max Rx Buffer Credit: 12
Max Rx Buffer Credit (ISL): 12
Default Rx Buffer Credit: 12
Default Rx Buffer Credit (ISL): 12
Default Rx Buffer Credit Multiplier: 0
Rx Buffer Credit change not allowed
Max Private Devices: 63
Hw Capabilities: 0xb
Connector Type: 0x0
FCOT info -
Min Speed: 1000
Max Speed: 2000
Module Type: 8
Connector Type: 7
Gigabit Eth Compliance Codes: 0
FC Transmitter Type: 3
Vendor Name: CISCO-AGILENT
Vendor ID: 0:48:211
Vendor Part Num: QFBR-5796L
Vendor Revision Level:
Trunk Info -
trunk vsans (allowed active) (1)

```

If the configuration of the switch port is configured as auto, and the point-to-point link still comes up as an FL port, verify that the HBA is configured as an NL port also.

Some HBAs support only NL mode. Verify the HBA capabilities with the vendor.

Interface UP and Connectivity Problems - Troubleshooting VSANs and Zones

If a server is not able to see a storage device, it may be because of a VSAN or zone misconfiguration.

Zone problems are more likely to happen than VSAN issues. This is because zone configuration and the overall zone protocols are more complex than VSAN configuration, and the VSAN membership can be verified using the CLI or the GUI. Therefore, checking the zone configuration is the first step to take when the host is not able to access the storage, and the port are all up along the path between the server HBA and the storage subsystem interface.

Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting Zones – Case of end devices belonging to the default zone

The first thing to check when performing zone troubleshooting is whether the storage subsystem port and the server HBAs have been configured to belong to a specific zone, or whether they belong to the default zone (any device that is not part of any active zone, it is considered to be part of the default zone).

In case the device does not belong to any zone, check whether the default zone default policy is set to “permit” on any switch in the fabric for the specific VSAN.

To set the default zone policy to “permit”, use the following command:

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

If the server is still not able to see the storage after you have configured the default-zone policy to “permit” on each switch in the fabric, check the VSAN configuration or the server and storage subsystem configuration.

Troubleshooting Zones – Case of end devices belonging to a specific zone

If zoning has been configured and the server and storage subsystem having the problem, it is important to check the correctness of the zoning configuration.

The following configuration steps should be followed in order to have zoning to work properly in the fabric or in a specific VSAN:

- The zone must be created in a VSAN.
- The correct FC IP, pwwn, or alias must be added to the Zone in the VSAN.
- A zoneset must be created in the VSAN.
- The zone must be added to the zoneset in the VSAN.
- The zone set must be activated in the VSAN.

It is important to verify the information contained in the active zoneset database for a particular VSAN. The active zoneset database is the only meaningful information for troubleshooting, because the active zoneset policy is applied to every switch in the fabric.

Verify active zoneset configuration

To verify that the zone has been created in a VSAN, use the **show zone active** command:

```
switch# show zone active
```

If there are configured active zones on the switch, the output of the command should look like:

```
switch# show zone active

zone name Zone1 vsan 1
  pwwn 50:06:0e:80:03:50:5c:03
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone2 vsan 1
  pwwn 10:00:00:e0:02:21:df:ef
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone3 vsan 1
  fwwn 20:42:00:05:30:00:63:9e
  fwwn 20:43:00:05:30:00:63:9e
```

Send comments to mdsfeedback-doc@cisco.com.

```
zone name zone-cc vsan 2
  pwwn 50:06:0e:80:03:50:5c:01
  pwwn 20:00:00:e0:69:41:a0:12
  pwwn 20:00:00:e0:69:41:98:93
```

If the output of this command doesn't show any information, this means that no zoneset has been activated in the fabric. If this is the case, you can use the two commands **show zone** or **show zoneset** to determine whether a zone configuration has been issued on the switch. (Zone configuration, and configuration changes do not get propagated to the other switches in the fabric, but only the changes to the active zones or active zoneset get propagated. For this reason if zone information is configured on one switch, it won't appear in the configuration of other switches).

Assuming that the zone configuration has been issued on the switch, but no active zoneset is shown, the next step is to enable the active zoneset in the specific VSAN to which the zoneset is supposed to belong.

To activate a zoneset in a defined VSAN, run the following command:

```
switch(config)# zoneset activate name ZonesetName vsan 1
```

The command must be issued on same switch the zone configuration took place. By copying the active zone database on the local zone database, it is possible to modify and apply those changes to the active zoneset on a different switch from the one initially used to issue the active zone configuration.

If no port shows as isolated, it means that all the switches in the fabric (or in the specific VSAN) share the same active zone database. This is ensured by the merge and change protocols used whenever any of the following occurs:

- two fabrics are connected
- a new ISL is configured in the fabric
- a new switch is connected to a pre-configured fabric
- changes are applied to the active zone database on any switch in the fabric

Verify active zoneset membership

In case no port shows as isolated, check that the HBA's FCID or pwwn and the storage subsystem FCID or pwwn are correctly configured to belong to the same zone. Or, if they have been added to an fcalias, check the correctness of the fcalias definition. Verify this information using the **show zone active** and **show zone** commands, or by using the Fabric Manager VSAN/Zoning View.

In the following example, only the server HBA's PWWN appears to be configured in the active zoneset belonging to zone-cc on VSAN 2.

```
switch# show zoneset active
zoneset name ZoneSet1 vsan 1
zone name Zon1 vsan 1
  pwwn 50:06:0e:80:03:50:5c:03
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone2 vsan 1
  pwwn 10:00:00:e0:02:21:df:ef
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone3 vsan 1
  fwwn 20:42:00:05:30:00:63:9e
  fwwn 20:43:00:05:30:00:63:9e

zone name zone-cc vsan 2
  pwwn 50:06:0e:80:03:50:5c:01
  pwwn 20:00:00:e0:69:41:a0:12
```

Send comments to mdsfeedback-doc@cisco.com.

```
pwwn 20:00:00:e0:69:41:98:93
```

To determine why, issue the **show zoneset** command.

```
switch# show zoneset
zoneset name ZoneSet1 vsan 1
zone name Zone1 vsan 1
  pwwn 50:06:0e:80:03:50:5c:03
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone2 vsan 1
  pwwn 10:00:00:e0:02:21:df:ef
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name Zone3 vsan 1
  fwwn 20:42:00:05:30:00:63:9e
  fwwn 20:43:00:05:30:00:63:9e

zone name zone-cc vsan 2
  pwwn 50:06:0e:80:03:50:5c:01
  pwwn 20:00:00:e0:69:41:a0:12
  pwwn 20:00:00:e0:69:41:98:93
zone name Zone-cc vsan 2
  pwwn 50:00:00:20:37:6f:db:aa.
```

The output shows that the zone name was added incorrectly (names for both zones and zoneset are case sensitive), creating a new zone called Zone-cc instead of zone-cc. Therefore, the storage pwwn has been erroneously configured to belong to a zone in which that port is the only member. The pwwn of the storage subsystem disappears from the output of **show active zoneset** command, because Zone-cc is not be added to the active zoneset of VSAN2.

Other useful commands

Use the **show zone name** command to display members of a specific zone.

```
switch# show zone name Zone1
zone name Zone1 vsan 1
```

Use the **show fcalias** to show if and how aliases are configured.

```
switch# show fcalias vsan 1
fcalias name Alias2 vsan 1 pwwn 21:00:00:20:37:6f:db:dd
fcalias name Alias1 vsan 1 pwwn 21:00:00:20:37:9c:48:e5
```

Use the **show zone member** command to display all zones to which a member belongs using the FCID, the fcalias, or the pwwn.

```
switch# show zone member pwwn 21:00:00:20:37:9c:48:e5
VSAN: 1
zone Zone3
zone Zone1
fcalias Alias1
```

Use the **show zone statistics** command to display the number of control frames exchanged with other switches.

```
switch# show zone statistics
```

Send comments to mdsfeedback-doc@cisco.com.

```
Statistics For VSAN: 1
*****
Number of Merge Requests Sent: 24
Number of Merge Requests Recvd: 25
Number of Merge Accepts Sent: 25
Number of Merge Accepts Recvd: 25
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
Statistics For VSAN: 2
*****
Number of Merge Requests Sent: 4
Number of Merge Requests Recvd: 4
Number of Merge Accepts Sent: 4
Number of Merge Accepts Recvd: 4
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
```

The **show zone internal vsan** command shows the internal state of the zone server for a specific VSAN.

```
switch# sh zone internal vsan 1
VSAN: 1 default-zone: deny distribute: active only
      E_D_TOV: 2000 R_A_TOV: 10000 F_S_TOV: 5000 Interop: Off
      DBLock:-(F count:0) Ifindex Table Size:2
Full Zoning Database :
      Zonesets:6 Zones:6 Aliases:0
Active Zoning Database :
      Name: ZoneSet6 Zonesets:1 Zones:1 Aliases:0
TCAM Info :
      cur_seq_num : 9, state : 0
      add_reqs = 4, del_reqs = 0, entries_added = 0
Change protocol info :
      local domain id = 102, ACA by 0xff
      State = Idle, reply_cnt = 0, req_pending = 0
      Remote domains :
Merge proto info :
      i/f fc2/15 | State = Isolated | notify = 0x8 | - -
```

Using the GUI to Troubleshoot Zoning Configuration Issues

Much of the information accessible through CLI commands can be accessed and summarized using the Fabric Manager VSAN/Zone view. For example, to check which devices belong to the active zoneset on a specific VSAN, click on the folder representing the active zoneset. This will display the set of devices belonging to that zoneset in that particular VSAN.

Similarly, by expanding the active zoneset folder content (clicking on the '+' next to the folder) the members of the active zoneset (the active zones) will be displayed as new folders.

Send comments to mdsfeedback-doc@cisco.com.

By recursively expanding the zone folders, the devices belonging to that zone will be listed in the left side column of the Fabric Manager window, and they will be highlighted in the topology view on the right side of the Fabric Manager window.

Troubleshooting VSANs

If VSANs are not configured properly, host devices will not be allowed to see storage devices configured to belong to different VSANs.

Hosts and storage ports must belong to the same VSAN, and VSANs cannot overlap.

VSAN membership for a specific port can be verified using the following command:

```
switch# show vsan membership interface fc2/1
fc2/1
      vsan:3
```

The output above shows interface fc2/1 is in VSAN 3.

To troubleshoot VSAN membership problems, issue the same command for both the port connected to the servers and the ones connecting the storage subsystem to the fabric. Then, verify that the VSAN is the same for both.

A more general command to verify the VSAN membership of all the ports on switch is:

```
switch# show vsan membership
vsan 1 interfaces:
      fc2/7  fc2/8  fc2/9  fc2/10  fc2/11  fc2/12  fc2/13  fc2/14
      fc2/15 fc2/16 fc7/1  fc7/2  fc7/3  fc7/4  fc7/5  fc7/6
      fc7/7  fc7/8  fc7/9  fc7/10  fc7/11  fc7/12  fc7/13  fc7/14
      fc7/15 fc7/16 fc7/17  fc7/18  fc7/19  fc7/20  fc7/21  fc7/22
      fc7/25 fc7/26 fc7/27  fc7/28  fc7/29  fc7/30  fc7/31  fc7/32

vsan 2 interfaces:
      fc2/6  fc7/23  fc7/24

vsan 3 interfaces:
      fc2/1  fc2/2  fc2/5

vsan 4 interfaces:
      fc2/3  fc2/4
```

If the devices are on different switches, issue the **show vsan membership** command on both devices. Then, verify that the trunks connecting the end switches are configured to transport the VSAN in question. This is done by issuing the **show interface** command, to verify that the port is in trunk mode and that the VSAN in question belongs to the trunk VSAN. Refer to the example below. If this is not the case, refer to the Troubleshooting ISL Isolation section at the beginning of this chapter to determine how to troubleshoot the connectivity issue.

```
switch# sh int fc2/14
fc2/14 is trunking
      Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e
      Port mode is TE
      Speed is 2 Gbps
      vsan is 2
      Beacon is turned off
      Trunk vsans (allowed active) (1-3,5)
      Trunk vsans (operational)    (1-3,5)
      Trunk vsans (up)            (2-3,5)
      Trunk vsans (isolated)      (1)
      Trunk vsans (initializing)  ()
      475 frames input, 8982 bytes, 0 discards
```

Send comments to mdsfeedback-doc@cisco.com.

```
0 runts, 0 jabber, 0 too long, 0 too short
0 input errors, 0 CRC, 3 invalid transmission words
0 address id, 0 delimiter
0 EOF abort, 0 fragmented, 0 unknown class
514 frames output, 7509 bytes, 16777216 discards
Received 30 OLS, 21 LRR, 18 NOS, 53 loop inits
Transmitted 68 OLS, 25 LRR, 28 NOS, 32 loop inits
```

Using the GUI to Troubleshoot VSAN Membership Problems

Checking the VSAN membership can also be done using the Device Manager.

Another tool that may be used to verify different categories of problems (VSANs, zoning, fcdomain, admin issues, or other switch-specific or fabric-specific issues) is the Fabric Analysis tool provided by Fabric Manager.

The possible configuration consistency check tool is also provided by this application. Refer to the *Cisco MDS 9000 Fabric Manager User Guide* for further information about this tool.

Troubleshooting Switch Fabric Level Issues

This chapter describes switch fabric-level troubleshooting procedures and includes the following sections:

- Troubleshooting Name Server Issues, page 4-1
- Troubleshooting FSPF Issues, page 4-6
- Troubleshooting Zoning Issues, page 4-14

Troubleshooting Name Server Issues

This section describes how to identify and resolve problems with the Cisco MDS 9000 Family name server. It includes the following sections:

- Overview, page 4-1
- Nx Port Registration Problems, page 4-2

Overview

The name server provides a way for N ports and NL ports to register and discover FibreChannel attributes. Registrations may happen explicitly, as a consequence of a request of the Nx port, or implicitly by the switch at FLOGI time. Once registered, the attributes are made available to other Nx ports or other switches. A separate name server database is maintained for each VSAN. The name server uses the Fibre Channel Common Transport protocol (FC-CT) to communicate with Nx ports and represents itself as an N port at the well-known FCID 0xFFFFFC.

One instance of the name server process runs on each Cisco MDS 9000 Family switch. In a multi-switch fabric configuration, instances running on different switches share information and create a distributed database of Nx port attributes. The name server defines a set of attribute objects with the following operations on those objects:

- Register Object
- Deregister Object
- Get Object

To troubleshoot name server problems, check the status of these three operations and verify the correct distribution of attributes among name server instances within the fabric. These attributes include the following:

- Port Type

Send comments to mdsfeedback-doc@cisco.com.

- Port Identifier
- Port Name
- Node Name
- Port Symbolic Name
- Node Symbolic Name
- Class of Service
- FC-4 Type(s)
- Port IP Addresses
- Node IP Addresses
- Hard Address
- FC-4 Descriptor
- FC-4 Features
- Initial Process Associator

Nx Port Registration Problems

To troubleshoot port registration, follow these steps:

Step 1 From the CLI exec mode, enter the following command:

```
show interface fcx/x
```

This ensures that the fibre channel (FC) interface connected to the device in question is up and free of any errors.

The system output might look like this:

```
switch# show int fc3/14
fc3/14 is up
  Hardware is Fibre Channel
  Port WWN is 20:8e:00:05:30:00:86:9e
  Admin port mode is FX
  Port mode is F, FCID is 0x780200 /* Operational State of the Port */
  Port vsan is 99 /* This is the vsan */
  Speed is 2 Gbps
  Receive B2B Credit is 16
  Receive data field size is 2112
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  1700 frames input, 106008 bytes, 0 discards
    0 CRC, 0 unknown class
    0 too long, 0 too short
  2904 frames output, 364744 bytes, 0 discards
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  1 output OLS, 1 LRR, 0 NOS, 0 loop inits
```

If the interface is not working correctly, check the cabling and the host or storage device interface for faults. If the interface is working correctly, proceed to the next step.

Step 2 Verify that the device in question appears in the FLOGI database. To do this, enter the following command:

Send comments to mdsfeedback-doc@cisco.com.

```
show flogi database vsan vsanid
```

The system output might look like this:

```
switch# show flogi database vsan 99
```

```
-----
INTERFACE  VSAN      FCID      PORT NAME      NODE NAME
-----
fc3/14      99        0x780200  21:00:00:e0:8b:07:a4:36  20:00:00:e0:8b:07:a4:36
```

If the device in question appears in this output, skip to Step 8. If the device does not appear in the output, go to the next step.

- Step 3** From interface mode, shut down the FC interface connected to the device in question.

```
config terminal
interface fcx/x
shutdown
```

- Step 4** Enter the following command on the FC interface:

```
no shutdown
```

By shutting down the interface and bringing it back up, you can determine what happens when the connected device tries to log in to the interface.

- Step 5** Enter the following command to view the events that occurred on the interface after you enabled it again:

```
switch# show flogi internal event-history interface fc3/14
```

The system output looks like this:

```
>>>>FSM: <[99]21:00:00:e0:8b:07:a4:36> has 9 logged transitions<<<<<
/* This is the [VSAN] followed by the pwn of the N/NL port */

1) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 321686 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_FLOGI_RECEIVED]
    Triggered event: [FLOGI_EV_VALID_FLOGI]
    Next state: [FLOGI_ST_GET_FCID]
/* The hba has sent an FLOGI to the switch */

2) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 322974 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_GET_FCID]
    Triggered event: [FLOGI_EV_VALID_FCID]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* Port Manager Obtains a valid FC_ID from the Domain Mgr */

3) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 323731 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_CONFIG_DONE_PENDING]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* ACLs are programmed and FIB {VSAN, FC_ID, portindex} is set */

4) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 323948 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_LCP_RESPONSE]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* LineCard responds that it is done */
```

Send comments to mdsfeedback-doc@cisco.com.

```

5) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 325962 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_NAME_SERVER_REG_RESPONSE]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* Program the NameServer with wwn and FCID */

6) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 330381 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_ZS_CFG_RESPONSE]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* Response from ZoneServer */

7) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 331187 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_RIB_RESPONSE]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* Response from RIB */

8) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 331768 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_ACL_CFG_RESPONSE]
    Next state: [FLOGI_ST_PERFORM_CONFIG]
/* Response from RIB */

9) FSM:<[99]21:00:00:e0:8b:07:a4:36> Transition at 331772 usecs after Sun Feb 1
04:18:15 1980
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_CONFIG_DONE_COMPLETE]
    Next state: [FLOGI_ST_FLOGI_DONE]
/* Programming done */

    Curr state: [FLOGI_ST_FLOGI_DONE]
/* Flogi was successful */

```

The comments that follow each section of output explain the meaning of the output.

If the device logs in successfully, proceed to the next step. Otherwise, you may have a problem with the device or its associated software.

Step 6 From interface mode, shut down the FC interface and issue a no shutdown after turning on the debug described in the following steps.

Step 7 To watch the FLOGI process take place, enter the following command:

```
switch# debug fcns events register vsan 99
```

This command enables debug mode for nameserver registration. The system output looks like this:

```

switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# int fc3/14
switch(config-if)# no shut /* enable the port */

switch(config-if)# Feb 17 04:42:54 fcns: vsan 99: Created entry for port-id 27800
Feb 17 04:42:54 fcns: vsan 99: Got Entry for port-id 27800
Feb 17 04:42:54 fcns: vsan 99: Registered port-name 36a4078be0000021 for port-id 780200
Feb 17 04:42:54 fcns: vsan 99: Registered node-name 36a4078be0000020 for port-id 780200
/* The wwpn and FCID for the port, note that the bytes in the world wide name are reversed
*/
Feb 17 04:42:54 fcns: vsan 99: Registered cos 8 for port-id 780200
/* Class of Service */

```

Send comments to mdsfeedback-doc@cisco.com.

```
Feb 17 04:42:54 fcns: vsan 99: Registered port-type 1 for port-id 780200
/* Port Type */
Feb 17 04:42:54 fcns: vsan 99: Reading configuration for entry with port-name
36a4078be0000021, node-name 36a4078be0000020
Feb 17 04:42:54 fcns: vsan 99: No configuration present for this portname
Feb 17 04:42:54 fcns: vsan 99: No configuration present for this nodename
/* Port is now registered in nameserver, will send out RSCN to it */

Feb 17 04:42:54 fcns: vsan 99: Trying to send RSCN; affected port 780200
Feb 17 04:42:54 fcns: vsan 99: rscn timer started for port 780200
Feb 17 04:42:54 fcns: vsan 99: Saving new entry into pss
Feb 17 04:42:54 fcns: vsan 99: Sending sync message to the standby
Feb 17 04:42:54 fcns: vsan 99: sending accept response to 780200
/* RSCN was received by N/NL port */

Feb 17 04:42:54 fcns: vsan 99: sending accept response to fffc61
/* Other switch in fabric is notified */
Feb 17 04:42:55 fcns: vsan 99: rscn timer expired for port 780200
Feb 17 04:42:55 fcns: vsan 99: Saving modified entry into pss
Feb 17 04:42:55 fcns: vsan 99: Sending sync message to the standby

Feb 17 04:42:55 fcns: vsan 99: Registered fc4-types for port-id 780200
Feb 17 04:42:55 fcns: vsan 99: Registered fc4-features for fc4_type 8 for port-id 780200
/* FC4 Type, type 8 FCP has been registered */
```

Additional lines like these will be listed if additional nameserver objects are registered

Step 8 From the CLI exec mode, enable FC name server (FCNS) debugging by entering the following command:

```
debug fcns events register vsan x
```

If you are managing the switch over a telnet connection, enable terminal monitoring by entering the **terminal monitor** command from the CLI exec mode.

The system output looks like this:

```
switch# show fcns database detail v 99
-----
VSAN:99      FCID:0x780200
-----
port-wwn (vendor)      :21:00:00:e0:8b:07:a4:36 (QLogic) /* Port world wide name */
node-wwn               :20:00:00:e0:8b:07:a4:36
class                 :3                               /* Fibrechannel class of service */
node-ip-addr          :0.0.0.0                         /* IP Address */
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init                    /* Registered FC4 Types: example SCSI and
initiator */
symbolic-port-name     :
symbolic-node-name     :
port-type              :N                               /* Fibrechannel port type (F,FL) */
port-ip-addr           :0.0.0.0
fabric-port-wwn        :20:8e:00:05:30:00:86:9e /* wwn of the switch port */
hard-addr              :0x000000
```

Other attribute objects of the Nx port are registered one per register operation after the FLOGI process is complete. The Nx port performs PLOGI to the well-known WWN of the Name Server, 0xFFFFFC. The FC_CT Common Transport protocol uses Request and Accept messages to conduct transactions. To verify that additional attributes are correctly registered and recorded in the database, you can use the SAN-OS debug facility.

Send comments to mdsfeedback-doc@cisco.com.



Note

The command **show fcns database detail vsan X** displays a detailed list of all devices registered in the fabric.

Troubleshooting FSPF Issues

This section describes how to identify and resolve FSPF problems and includes the following sections;

- Overview, page 4-6
- Loss of Two-way Communication, page 4-7
- Wrong Hello Interval on an ISL Interface, page 4-7
- Resolving the Wrong Hello Interval Problem, page 4-8
- Wrong Dead Interval on an ISL Interface, page 4-8
- Resolving a Wrong Dead Interval Problem, page 4-9
- Region Mismatch on Switch, page 4-9
- Resolving a Region Mismatch Problem, page 4-10
- FSPF Issues in a Single-VSAN Environment, page 4-11
- FSPF Issues in a Multi-VSAN Environment, page 4-13

Overview

To see all the correct FSPF information, as shown in the previous examples, the switches must be configured correctly. If FSPF is misconfigured, then the switches will not reach the “two-way” state. This can happen when:

- The switch fails to receive Hello in expected time (dead interval)
- The switch receives Hello from neighbor that does not contain the correct domain ID in the Recipient Domain ID field.



Note

This would not be a configuration issue.

- The switch receives Hello from neighbor with 0xFFFFFFFF in the Recipient Domain ID field.



Note

Hellos are sent with 0xFFFFFFFF in neighbor field until switch learns its neighbor’s domain ID.

- The switch receives Hello with incorrect Hello and dead intervals.

Send comments to mdsfeedback-doc@cisco.com.

Loss of Two-way Communication

The following events occur when two-way communication is lost:

1. The Port enters Init state and removes its neighbor's domain ID from the Recipient Domain ID field and inserts 0xFFFFFFFF.
2. FSPF Removes the ISL (inter switch link) from the topology database.
3. New LSRs (Link State Records) are flooded to adjacent switches to notify them that the FSPF database has changed.

Wrong Hello Interval on an ISL Interface

To identify a mismatch in the Hello interval on the two sides of an ISL interface, follow these steps:

Step 1 To turn on debugging, enter the following command:

```
switch1# debug fspf all
```

The system output looks like this:

```
switch1# debug fspf all
Jan 5 00:28:14 fspf: Wrong hello interval for packet on interface 100f000 in VSAN 1
Jan 5 00:28:14 fspf: Error in processing hello packet , error code = 4
```

Step 2 To show FSPF information, enter the following command:

```
switch1# show fspf internal route v 1
```

The system output looks like this:

```
switch1# show fspf internal route v 1
FSPF Unicast Routes
-----
VSAN Number      Dest Domain    Route Cost    Next hops
-----
1                0xEF(239)     1000          fc1/1 -----1
1                0xED(238)     2000          fc1/1
1                0x01(1)       3000          fc1/1 -----2
```

1. Indicates that there is no second path to Domain 238, through Domain 1 switch2.
2. Indicates that there is no direct path to Domain 1 switch2; traffic must travel through 3 ISLs.

Step 3 To view the currently configured FSPF parameters on the ISL, enter the following command:

```
switch1# show fspf v 1 interfac fc1/16
```

The system output looks like this:

```
switch1# show fspf v 1 interfac fc1/16
FSPF interface fc1/16 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 5 s, Dead 80 s, Retransmit 5 s -----1
FSPF State is INIT -----2
Statistics counters :
  Number of packets received : LSU 0 LSA 0 Hello 2 Error packets 1
  Number of packets transmitted : LSU 0 LSA 0 Hello 4 Retransmitted LSU 0
Number of times inactivity timer expired for the interface = 0
```

Send comments to mdsfeedback-doc@cisco.com.

1. Shows that the hello timer is not set to the default, so you should check the neighbor configuration to make sure it matches.
2. Shows that FSPF is not in FULL state, indicating a problem.

Step 4 To determine the value of the Hello timer on the adjacent switch, enter the following command:

```
NEIGHBOR
switch2# show fspf v 1 interfac fc1/16
```

The system output looks like this:

```
no shutdown
NEIGHBOR
switch2# show fspf v 1 interfac fc1/16
FSPF interface fc1/16 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s -----1
FSPF State is INIT -----2
Statistics counters :
  Number of packets received : LSU 0 LSA 0 Hello 2 Error packets 1
  Number of packets transmitted : LSU 0 LSA 0 Hello 4 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

1. Shows that the neighbor FSPF Hello interval is set to the default (20 seconds).
2. Indicates that FSPF is not in FULL state, indicating a problem.

Step 5 An alternative to check the hello interval setting is the show run command.

```
switch1# show run
interface fc1/1
 fspf hello-interval 5 vsan 1
no shutdown
```

The system output looks like this.

```
switch1# show run
interface fc1/1
no shutdown
```

This output indicates that the neighbor FSPF hello is set to the default. The default setting does not display in the output from the **show run** command.

Resolving the Wrong Hello Interval Problem

The Hello interval must match on both sides of an ISL, so you should change either side to match the other. Use the default Hello interval unless you are sure you need change it. To change the default Hello interval, enter the following commands:

```
ML-88(config)# interface fc 1/16
ML-88(config-if)# fspf hello-interval XX vsan 1
```

Wrong Dead Interval on an ISL Interface

To identify a mismatch in dead intervals on the two sides of an ISL interface, follow these steps:

Send comments to mdsfeedback-doc@cisco.com.

- Step 1** To enable the debug output for identifying this condition, enter the following command:

```
switch1# debug fspf all
```

The system output looks like this:

```
switch1# debug fspf all
Jan  5 00:28:14 fspf: Wrong dead interval for packet on interface 100f000 in VSAN 1
Jan  5 00:28:14 fspf: Error in processing hello packet , error code = 4
```

- Step 2** To display the currently configured FSPF parameters on the interface, enter the following command:

```
switch1# show fspf v 1 interfac fcl/16
```

You should run this command on the local interface and on the switch at the other end of the ISL on which the problem is occurring.

The system output looks like this:

```
switch1# show fspf v 1 interfac fcl/16
FSPF interface fcl/16 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 95 s, Retransmit 5 s -----1
FSPF State is INIT -----2
XStatistics counters :
  Number of packets received : LSU 0 LSA 0 Hello 2 Error packets 1
  Number of packets transmitted : LSU 0 LSA 0 Hello 4 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

1. Indicates that the dead timer is not set to the default, so you should check the neighbor configuration.
2. Indicates that FSPF is not in full state, which indicates a problem.

Resolving a Wrong Dead Interval Problem

The dead interval must match on both sides of the ISL, so to resolve this problem, change either side to match the other. Note that under most conditions, the default dead interval value should be used unless there is a demonstrated need to change it.

```
ML-88(config)# interface fc 1/16
ML-88(config-if)# fspf dead-interval XX vsan 1
```

Region Mismatch on Switch

To identify a region mismatch problem on a switch, follow these steps:

- Step 1** To display the currently configured region in a VSAN, enter the following command:

```
switch# show fspf vsan 99
```

The system output looks like this:

```
FSPF routing for VSAN 99
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0 /* This is the region */
```

Send comments to mdsfeedback-doc@cisco.com.

```
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x78(120)
Number of LSRs = 2, Total Checksum = 0x000133de
```

Step 2 To turn on debugging, enter the following command:

```
switch1# debug fspf all
```

The system output looks like this:

```
switch1# debug fspf all
Jan 5 00:39:31 fspf: FC2 packet received for non existent region 0 in VSAN 1 -----1
Jan 5 00:39:33 fspf: FC2 packet received for non existent region 0 in VSAN 1
Jan 5 00:39:45 fspf: Interface fcl/1 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT
Jan 5 00:39:45 fspf: Interface fcl/2 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT -----2
```

1. Indicates that the neighbor switch advertising region is 0.
2. Indicates that FSPF is in INIT state for each ISL.

Step 3 To check region settings, issue a show run command.

```
switch1# show run
fspf config vsan 1
    region 1
```

Resolving a Region Mismatch Problem

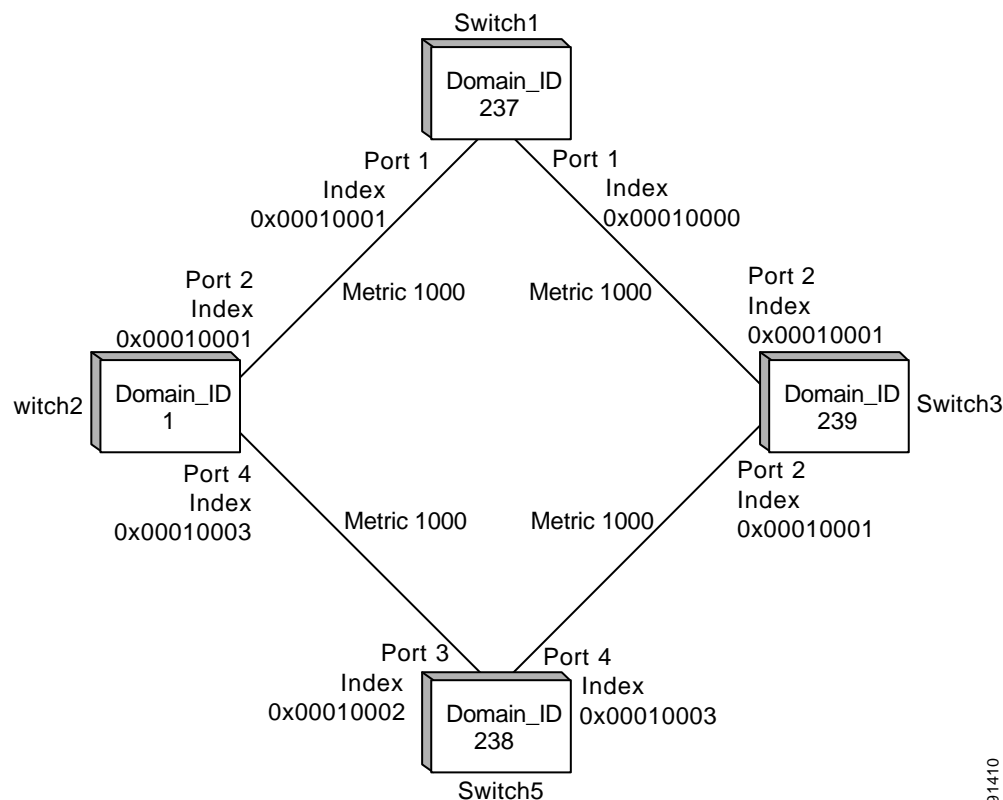
The region must match on all switches in the VSAN. To correct a region mismatch, enter the following commands:

```
switch1(config)# fspf config vsan 1
switch1(config-(fspf-config))# region 0
```

Send comments to mdsfeedback-doc@cisco.com.

FSPF Issues in a Single-VSAN Environment

Figure 4-1 Single VSAN Topology



91410

For the purpose of this example, assume that all interfaces are located in VSAN 1.

- Step 1** Verify that each path is in the FSPF database by entering the following command at the exec prompt:

```
Switch switch1# show fspf database
```

The system output looks like this:

```
Switch switch1# show fspf database
```

```
FSPF Link State Database for VSAN 1 Domain 1 -----1
```

```
LSR Type                = 1
Advertising domain ID    = 1 -----2
LSR Age                  = 81 -----3
LSR Incarnation number   = 0x80000098 -----4
LSR Checksum             = 0x2cd3
Number of links          = 2
```

NbrDomainId	IfIndex	NbrIfIndex	Link Type	Cost
237	0x00010002	0x00010001	1	1000 -----5
238	0x00010003	0x00010002	1	1000 -----6

The following is the beginning of another switch's LSR (Link State Record).

```
FSPF Link State Database for VSAN 1 Domain 237
```

```
LSR Type                = 1
Advertising domain ID    = 237 -----7
```

Send comments to mdsfeedback-doc@cisco.com.

```

LSR Age                = 185
LSR Incarnation number = 0x8000000c
LSR Checksum           = 0xe0a2
Number of links        = 2
  NbrDomainId          IfIndex      NbrIfIndex      Link Type      Cost
-----
239                    0x00010000      0x00010003      1              1000 -----8
1                      0x00010001      0x00010002      1              1000 -----9

```

The following is the beginning of another switch's LSR (Link State Record)

```

FSPF Link State Database for VSAN 1 Domain 238
LSR Type                = 1
Advertising domain ID   = 238
LSR Age                 = 1052
LSR Incarnation number = 0x80000013
LSR Checksum            = 0xe294
Number of links         = 2
  NbrDomainId          IfIndex      NbrIfIndex      Link Type      Cost
-----
239                    0x00010003      0x00010001      1              1000
1                      0x00010002      0x00010003      1              1000

```

The following is the beginning of another switch's LSR (Link State Record)

```

FSPF Link State Database for VSAN 1 Domain 239
LSR Type                = 1
Advertising domain ID   = 239
LSR Age                 = 1061
LSR Incarnation number = 0x80000086
LSR Checksum            = 0x66ac
Number of links         = 4
  NbrDomainId          IfIndex      NbrIfIndex      Link Type      Cost
-----
237                    0x00010003      0x00010000      1              1000
238                    0x00010001      0x00010003      1              1000

```

1. Provides the Domain 1 view of the fabric topology.
2. Indicates that Domain 1 is owner of the LSR (Link State Record).
3. This is a 16-bit counter starting at 0x0000, incremented by one for each switch during flooding and by one for each second held in database. This field is used as a tie-breaker if Incarnation numbers are the same.
4. This is a 32-bit value between 0x80000001 and 0x7FFFFFFF, which is incremented by one each time the originating switch transmits an LSR. This is used first before LSR Age.
5. Indicates the path to Domain 237, Switch switch1.
6. Indicates the path to Domain 238, Switch switch5.
7. Indicates that switch1, Domain ID 237 is the owner.
8. Indicates the path to Domain 239, Switch switch3.
9. Indicates the path to Domain 1, Switch switch2

Step 2 Verify that the FSPF parameters are correct for each interface by entering the following command at the exec prompt:

```
switch1# show fspf vsan 1 interface fc1/2
```

Send comments to mdsfeedback-doc@cisco.com.

View the output from this command to verify that the interface is in FSPF “active state.” The system output looks like this:

```
switch1# show fspf vsan 1 interface fc1/2
FSPF interface fc1/2 in VSAN 1
FSPF routing administrative state is active -----1
Interface cost is 1000 -----2
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s -----3
FSPF State is FULL -----4
Neighbor Domain Id is 1, Neighbor Interface index is 0x00010002 -----5
Statistics counters :
    Number of packets received : LSU 46 LSA 24 Hello 103 Error packets 0
    Number of packets transmitted : LSU 24 LSA 45 Hello 104 Retransmitted LSU 0
    Number of times inactivity timer expired for the interface = 0
```

This displays the number of packets; Hellos should be received every 20 seconds.

1. Indicates that FSPF is active and is not disabled on this interface.
2. Indicates the cost of the path out this interface.
3. Identifies the configured FSPF timers for this interface, which must match on both sides.
4. Indicates Full State or Adjacent. Sent and received all database exchanges and required Acks. Port is now ready to route frames.
5. Provides FSPF neighbor information.

Step 3 Verify that all FC routes are available by entering the following command:

```
switch1# show fspf internal route v 1
```

The system output looks like this:

```
switch1# show fspf internal route v 1
FSPF Unicast Routes
-----
  VSAN      Number      Dest Domain      Route Cost      Next hops
-----
  1          0x01(1)      1000             fc1/2
  1          0xEF(239)   1000             fc1/1
  1          0xED(238)   2000             fc1/1
                                     fc1/2
```

This shows the total cost of all links.

fc1/2

The next hop to (238) has two interfaces. This indicates that both paths will be used during load sharing. Up to sixteen paths can be used by FSPF with a Cisco MDS 9000 Family switch.

FSPF Issues in a Multi-VSAN Environment

With the implementation of VSANs used with Cisco MDS 9000 Family switches, a separate instance of FSPF runs within each VSAN, and each instance is independent of the others. For this reason, FSPF issues affecting one VSAN have no effect on FSPF running in other VSANs.

Send comments to mdsfeedback-doc@cisco.com.



Note

For all FSPF configuration statements and diagnostic commands, if the **vsan** keyword is not specified, VSAN 1 is used by default. When making configuration changes or issuing diagnostic commands in a multi-VSAN environment, be sure to explicitly specify the target VSAN by including the **vsan** keyword in the statement or command.

Troubleshooting Zoning Issues

This section describes how to identify and resolve zoning issues that may arise in a multiswitch fabric. It includes the following topics:


- Mismatched Active Zonesets Within the Same VSAN, page 4-14
- Importing or Exporting a Zoneset Between Switches, page 4-16
- Deactivating a Zoneset and Restarting the Zone Merge Process, page 4-17
- Misconfigured Zones Within an Active Zoneset in the Same VSAN, page 4-19

If after you verify the proper operation of the fibre channel name server and FSPF problems remain discovering remote switches and their attached resources, the fabric may have zone configuration problems. Examples of zone configuration problems are mismatched active zonesets and misconfigured zones within the active zoneset.

Mismatched Active Zonesets Within the Same VSAN

When merging switch fabrics, you must ensure that the zones in both active zonesets have unique names, or that any zones with the same name have exactly the same members. If either of these conditions is violated the E port connecting the two fabrics will appear in an isolated state.

Figure 4-2 Topology for Zone Merge Failure Example



In this example, two switches have the same zoneset name, and the same zone names, but different zone members. As a result, the VSAN is isolated on the TE port that connects to two switches.

This issue can be resolved by doing one of the following:

1. Modify the zone members on both zonesets to match and eliminate conflict
2. Deactivate the zoneset on one of the switches and restart the zone merge process
3. Explicitly import or export a zoneset between the switches to synchronize them.

To identify this problem, follow these steps:

Step 1 Enter the following command to display the active zoneset configuration of the first switch:

```
Switch1# show zoneset active v 99
```


Send comments to mdsfeedback-doc@cisco.com.

The system output looks like this:

```
Switch1# show zoneset active v 99
zoneset name ZoneSet1 vsan 99
  zone name VZ1 vsan 99
    * fcid 0x7800e2 [pwwn 22:00:00:20:37:04:ea:2b]
    * fcid 0x7800d9 [pwwn 22:00:00:20:37:04:f8:a1]
```

Step 2 Enter the following command to display the active zoneset configuration of the second switch:

```
Switch2# show zoneset active v 99
```

The system output looks like this:

```
Switch2# show zoneset active v 99
zoneset name ZoneSet1 vsan 99
  zone name VZ1 vsan 99
    pwwn 22:00:00:20:37:04:f8:a1
    pwwn 22:00:00:20:37:0e:65:44
```

Even though the zones have the same name, their respective members are different.

Step 3 Enter the following command to view information about the TE port:

```
Switch2# show int fc1/8
```

This command shows all the information about the interface. The system output looks like this:

```
Switch2# show int fc1/8
fc1/8 is trunking
  Hardware is Fibre Channel
  Port WWN is 20:08:00:05:30:00:5f:1e
  Peer port WWN is 20:05:00:05:30:00:86:9e
  Admin port mode is E, trunk mode is auto
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Receive data field size is 2112
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1,99)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (99)
  Trunk vsans (initializing) ()
  5 minutes input rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
  5 minutes output rate 88 bits/sec, 11 bytes/sec, 0 frames/sec
    10845 frames input, 620268 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    10842 frames output, 487544 bytes, 0 discards
      3 input OLS, 4 LRR, 3 NOS, 0 loop inits
      18 output OLS, 2 LRR, 14 NOS, 0 loop inits
```

From this output, you can see that VSAN 99 is isolated.

Step 4 Enter the following command to get information about why the interface is isolated:

```
switch2# show int fc1/8 trunk vsan
```

The system output looks like this:

```
switch2# show int fc1/8 trunk vsan
fc1/8 is trunking
  Vsan 1 is up, FCID is 0x650000
  Vsan 99 is down
```

Send comments to mdsfeedback-doc@cisco.com.

From this output, you can see the VSAN is isolated due to a zone merge failure:

Step 5 To resolve the isolation problem, do one of the following:

- Change the membership of one of the zones to match the other of the same name.
- Discard one of the zonesets completely by either deactivating it, or by overwriting the active zoneset on one switch using the **import** or **export** commands. This method is destructive to one of the active zonesets.

For instructions about changing the membership of a zone, refer to the *Cisco MDS 9000 Family Configuration Guide*.

Step 6 The following example shows the use of the **export** command:

In this example, the zoneset configuration is correct on switch1, so you would want to discard the zoneset configuration on switch2. You get the same result by entering the **zone import** command from switch2.

Step 7 Import the active zoneset for VSAN 99 on interface fc1/8 by entering the following command:

```
Switch1# zone merge interface fc1/8 import vsan 99
Zoneset export initiated. check zone status
```

Step 8 Verify that VSAN 99 is no longer isolated by entering the following command:

```
Switch1# show int fc1/5 trunk v 99
fc1/5 is trunking
    Vsan 99 is up, FCID is 0x780102
```

If a VSAN does not have an active zoneset, it automatically takes the active zoneset of the other merging switch. So another way to solve this problem is by deactivating the active zoneset on switch2 by entering the following command:

```
no zoneset activate name ZoneSet1 v 99
```

This removes the active zoneset on switch2, which will then automatically take the active zoneset from switch1.

Importing or Exporting a Zoneset Between Switches

To import or export a zoneset between switches, follow these steps:

Step 1 Enter the following command <<explain why?>>

```
switch4#
```

The system output looks like this:

```
Nov 19 09:28:45 switch4 %LOG_PORT_CHANNEL-5-FOP_CHANGED: port-channel 1: first operational
port changed from none to fc1/14
Nov 19 09:28:55 switch4 %LOG_ZONE-2-ZS_MERGE_FAILED: Zone merge failure, Isolating port
port-channel 1 (VSAN 1)
```

Step 2 Enter the following command <<explain why?>>

```
switch4# zone merge int port-channel 1 import vsan 1
```

The system output looks like this:

```
Zoneset Import initiated. check zone status
```

Send comments to mdsfeedback-doc@cisco.com.

```
switch4# Nov 19 11:43:02 switch4 %LOG_PORT-5-IF_UP: Interface port-channel 1 is up in mode
TE
Nov 19 11:43:02 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/14, vsan 1 is up
Nov 19 11:43:02 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/15, vsan 1 is up
Nov 19 11:43:02 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/16, vsan 1 is up
switch4# show zoneset activ
```

The system output looks like this:

```
zoneset name wall vsan 1
  zone name excall vsan 1
    * fcid 0x620200
      fcid 0x6200ca
zone name $default_zone$ vsan 1
  * fcid 0x6e00da
  * fcid 0x6e00d9
  * fcid 0x6e00d6
  * fcid 0x6201d5
  * fcid 0x6201d4
  * fcid 0x6201d3
  * fcid 0x6201d2
  * fcid 0x6201d1
  * fcid 0x6201ce
  * fcid 0x6201cd
  * fcid 0x6201cc
  * fcid 0x6201cb
  * fcid 0x6201ca
  * fcid 0x6e0100
switch4#
```

When you explicitly import a zoneset for VSAN 1 from switch3 on switch4 over the isolated ISL, the active zoneset is copied from switch3 to switch4. The zone databases are now synchronized, and the VSAN is no longer isolated.

Deactivating a Zoneset and Restarting the Zone Merge Process

To deactivate a zoneset and restart the zone merge process, follow these steps:

-
- Step 1** Remove the zoneset configuration from the switch, as in the following example:

```
switch4(config)# no zoneset activate name excal2 vsan 1
```

The system output looks like this:

```
Zoneset Deactivation initiated. check zone status
```

- Step 2** To confirm that the zoneset has been removed, enter the following command:

```
switch4(config)# exit
switch4# show zoneset active
```

- Step 3** Shut down the connection to the zone to be merged.

```
switch4# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch4(config)# int port-channel 1
switch4(config-if)# shut
```

The system output looks like this:

```
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/14 is down
(Channel admin down)
```

Send comments to mdsfeedback-doc@cisco.com.

```
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fcl/15 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fcl/16 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface
port-channel 1 is down (No operational members)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_ADMIN_DOWN: Interface port-channel 1 is down
(Administratively down)
Nov 19 10:26:10 switch4 %LOG_PORT_CHANNEL-5-FOP_CHANGED: port-channel 1: first operational
port changed from fcl/16 to none
```

Step 4 To reactivate the connection to the zone to be merged, enter the following command:

```
switch4(config-if)#
switch4(config-if)# no shut
```

The system output looks like this:

```
switch4(config-if)# Nov 19 10:28:11 switch4 %LOG_PORT_CHANNEL-5-FOP_CHAN
GED: port-channel 1: first operational port changed from none to fcl/15
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_UP: Interface port-channel 1 is up in mode TE
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/16, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fcl/16, vsan 1 is up
```

Step 5 To exit config mode and check the active zone sets, enter the following command:

```
switch4(config-if)# exit
switch4(config)# exit
switch4# show zoneset active
```

The system output looks like this:

```
zoneset name wall vsan 1
  zone name excall vsan 1
    * fcid 0x620200
    * fcid 0x6200ca
zone name $default_zone$ vsan 1
  * fcid 0x6e00da
  * fcid 0x6e00d9
  * fcid 0x6e00d6
  * fcid 0x6e0100
  * fcid 0x6201ca
  * fcid 0x6201cb
  * fcid 0x6201cc
  * fcid 0x6201cd
  * fcid 0x6201ce
  * fcid 0x6201d1
  * fcid 0x6201d2
  * fcid 0x6201d3
  * fcid 0x6201d4
  * fcid 0x6201d5
switch4#
```

After deactivating the zoneset on switch4 and performing a shutdown followed by a no shutdown on the ISL that connects it to switch3, the zone merge is processed again. Because switch3 has no active zoneset, it learns the active zoneset from switch4 during the zone merge process.

Send comments to mdsfeedback-doc@cisco.com.

Misconfigured Zones Within an Active Zoneset in the Same VSAN

Even when the active zonesets contain the same zones for a VSAN on all the switches within a fabric, the members contained within those zones must also match or the zone merge will fail.

In this example, the active zonesets “wall” for VSAN 1 on switches switch3 and switch4 contain the same zone “excal1”. However, the members of the zone are different on each switch. When the switches they are connected together to form the switch fabric, a zone merge failure occurs and VSAN 1 becomes isolated.

This topic has already been covered in a previous section. All you need to do is to modify the membership of the zones so that any zones that have the same names also have the exact same membership.

For instructions about changing the membership of a zone, refer to the *Cisco MDS 9000 Family Configuration Guide*.

Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting IP Storage Issues

This chapter describes IP storage troubleshooting procedures and includes the following sections:

- Overview, page 5-21
- Troubleshooting IP Connections, page 5-22
- Troubleshooting FCIP Connections, page 5-25
- Troubleshooting iSCSI Issues, page 5-51
- Fine Tuning/Troubleshooting IPS iSCSI TCP Performance, page 5-64


Overview

Using open-standard, IP-based technology, the Cisco MDS 9000 Family IP storage module enables you to extend the reach of Fibre Channel SANs. The switch can connect separated SAN islands together via IP networks using FCIP, and allow IP hosts to access FC storage using the iSCSI protocol.

The IP Storage (IPS) services module allows you to use FCIP and iSCSI features. It supports the full range of features available on other switching modules, including VSANs, security, and traffic management. The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports. Each port can run the FCIP and iSCSI protocols simultaneously.

FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices (see Figure 5-1). Using the iSCSI protocol, the IPS module provides IP hosts access to Fibre Channel storage devices. IP host-initiated iSCSI commands are encapsulated in IP, and sent to an MDS 9000 IPS port. There, the commands are routed from the IP network into a Fibre Channel network, and forwarded to the intended target.

Figure 5-1 Connecting MDS 9000 Family Switches Over IP



Send comments to mdsfeedback-doc@cisco.com.

Troubleshooting IP Connections

If you suspect that all or part of your IP connection has failed, you can verify that by performing one or more of the procedures in this section. Using these procedures, you can verify connectivity for IP, subinterfaces/802.1q, EtherChannel, and VRRP for iSCSI.

Verifying Basic Connectivity

- Step 1** Perform a basic check of host reachability and network connectivity using the **ping** command. A sample output of the **ping** command follows:

```
switch# ping 172.18.185.121
PING 172.18.185.121 (172.18.185.121): 56 data bytes
64 bytes from 172.18.185.121: icmp_seq=0 ttl=128 time=0.3 ms
64 bytes from 172.18.185.121: icmp_seq=1 ttl=128 time=0.1 ms
64 bytes from 172.18.185.121: icmp_seq=2 ttl=128 time=0.2 ms
64 bytes from 172.18.185.121: icmp_seq=3 ttl=128 time=0.2 ms
64 bytes from 172.18.185.121: icmp_seq=4 ttl=128 time=0.1 ms
64 bytes from 172.18.185.121: icmp_seq=5 ttl=128 time=0.1 ms

--- 172.18.185.121 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.3 ms
```

- Step 2** Verify route to remote device using **sh ip route**, **traceroute**, and **sh arp** commands. A sample output of the **sh ip route** command follows:

```
switch # sh ip route

Codes: C - connected, S - static

Default gateway is 172.18.185.97

C 172.18.185.96/27 is directly connected, mgmt0
C 172.18.189.128/26 is directly connected, gigabitethernet4/7
```

A sample output of the **traceroute** command follows. The route is using int GigE, verified using the **sh arp** command.

```
switch# traceroute 172.18.185.121
traceroute to 172.18.185.121 (172.18.185.121), 30 hops max, 38 byte packets
 1 172.18.185.121 (172.18.185.121) 0.411 ms 0.150 ms 0.146 ms
```

Another sample output of the **traceroute** command follows. This route is using int mgmt0, verified using the **sh arp** command.

```
switch# traceroute 10.82.241.17
traceroute to 10.82.241.17 (10.82.241.17), 30 hops max, 38 byte packets
 1 172.18.189.129 (172.18.189.129) 0.413 ms 0.257 ms 0.249 ms
 2 172.18.0.33 (172.18.0.33) 0.296 ms 0.260 ms 0.258 ms
 3 10.81.254.69 (10.81.254.69) 0.300 ms 0.273 ms 0.277 ms
 4 10.81.254.118 (10.81.254.118) 0.412 ms 0.292 ms 0.287 ms
 5 10.83.255.81 (10.83.255.81) 0.320 ms 0.301 ms 0.310 ms
 6 10.83.255.163 (10.83.255.163) 0.314 ms 0.295 ms 0.279 ms
 7 10.82.241.17 (10.82.241.17) 48.152 ms 48.608 ms 48.423 ms
```


Send comments to mdsfeedback-doc@cisco.com.

A sample output of the **sh ips arp** command follows.

```
switch# sh ips arp int giga 4/7
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet  172.18.185.97      0    00:d0:01:3b:38:0a  ARPA   GigabitEthernet4/7
Internet  172.18.189.129     0    00:d0:01:3b:38:0a  ARPA   GigabitEthernet4/7
Internet  172.18.189.153     0    00:08:02:24:e0:8b  ARPA   GigabitEthernet4/7
Internet  172.18.189.155     0    00:08:02:df:93:77  ARPA   GigabitEthernet4/7
Internet  172.18.189.156     9    00:08:02:b3:45:1b  ARPA   GigabitEthernet4/7
```

A sample output of the **clear ips arp** command follows. You clear the arp cache to verify that the activity you are viewing is the most current.

```
switch# clear ips arp int gig4/7
arp clear successful
```

A sample output of the **sh ips arp** command follows.

```
switch# sh ips arp int giga 4/7
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet  172.18.185.97      0    00:d0:01:3b:38:0a  ARPA   GigabitEthernet4/7
Internet  172.18.189.156     0    00:08:02:b3:45:1b  ARPA   GigabitEthernet4/7
```

A sample output of the **sh ips arp** command follows.

```
switch# sh ips arp int giga 4/7
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet  172.18.185.97      0    00:d0:01:3b:38:0a  ARPA   GigabitEthernet4/7
Internet  172.18.189.129     0    00:d0:01:3b:38:0a  ARPA   GigabitEthernet4/7
Internet  172.18.189.156     0    00:08:02:b3:45:1b  ARPA   GigabitEthernet4/7
```

A sample output of the **sh arp** command follows.

```
switch# sh arp
Protocol  Address      Age (min)    Hardware Addr  Type   Interface
Internet  172.18.185.97      0    00d0.013b.380a  ARPA   mgmt0
```

Step 3 Use the **show interface** command to verify that the Gigabit Ethernet interface is up. A sample output of the **show interface** command follows.

```
GigabitEthernet4/7 is up
  Hardware is GigabitEthernet, address is 0005.3000.9f58
  Internet address is 172.18.189.137/26
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 688 bits/sec, 86 bytes/sec, 0 frames/sec
  5 minutes output rate 312 bits/sec, 39 bytes/sec, 0 frames/sec
  156643 packets input, 16859832 bytes
    0 multicast frames, 0 compressed
```

Send comments to mdsfeedback-doc@cisco.com.

```
0 input errors, 0 frame, 0 overrun 0 fifo
144401 packets output, 7805631 bytes, 0 underruns
0 output errors, 0 collisions, 0 fifo
0 carrier errors
```

Verifying Static IP Routing

Static routing is a mechanism to configure IP routes on the switch. To verify that the IP routes are still there, use the **sh ip route** command. A sample output of the **sh ip route** command follows.

```
switch# sh ip route

Codes: C - connected, S - static

Default gateway is 172.17.8.1

C 172.17.8.0/24 is directly connected, mgmt0
S 11.2.36.0/22 via 11.3.36.1, gigabitethernet8/7
C 11.3.36.0/22 is directly connected, gigabitethernet8/7
C 11.3.56.0/22 is directly connected, gigabitethernet8/8
S 11.2.56.0/22 via 11.3.56.1, gigabitethernet8/8
switch#
```

A sample output of the **sh ip route config** command follows.

```
switch# sh ip route config
```

Destination	Gateway	Mask	Metric	Interface
default	172.17.8.1	0.0.0.0	0	mgmt0
11.2.36.0	11.3.36.1	255.255.252.0	0	
11.2.56.0	11.3.56.1	255.255.252.0	0	
11.3.36.0	0.0.0.0	255.255.252.0	0	GigabitEthernet8/7
11.3.56.0	0.0.0.0	255.255.252.0	0	GigabitEthernet8/8
172.17.8.0	0.0.0.0	255.255.255.0	0	mgmt0

```
switch#
```

Send comments to mdsfeedback-doc@cisco.com.


Troubleshooting FCIP Connections

This section contains information on troubleshooting FCIP tunnels with and without Special Frames.

One-to-One FCIP Tunnel Creation and Monitoring

This section describes the configuration for one-to-one FCIP tunnel with FCIP debug activated (MDS2) and without debug activated (MDS1). Figure 5-2 shows the one-to-one topology used for configuration.

Figure 5-2 One-to-One Topology



First, perform the following steps to configure the MDS1.

-
- Step 1** Enter configuration mode
- Step 2** Set the interface
- ```
MDS1(config)#interface gig 2/8
```
- Step 3** Set the IP address
- ```
MDS1(config-if)#ip address 10.10.10.2 255.255.255.0
```
- Step 4** Enter no shutdown for some reason
- ```
MDS1(config-if)#no shutdown
```
- Step 5** Enter the profile number and profile mode.
- ```
MDS1(config)# fcip profile 28
```
- The profile number can be any number between 1 – 255
- Step 6** Enter the IP address of the local GE port that will be endpoint of FCIP tunnel.
- ```
MDS1(config-profile)# ip address 10.10.10.2
```
- Step 7** Exit profile mode.
- ```
MDS1(config-profile)# exit
```
- Step 8** Set the interface FCIP and enter interface mode.
- ```
MDS1(config)# interface fcip 28
```

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

The interface FCIP can be any number between 1 – 255 and does not need to be the same as the profile number. In this example the same number is used for simplicity.

**Step 9** Specify a profile to use.

```
MDS1(config-if)# use-profile 28
```

The interface FCIP will use the Local FCIP profile . The FCIP profile binds the interface FCIP to the physical Gigabit Ethernet port and configures the TCP settings used by the interface FCIP.

```
MDS1(config-if)# peer-info ipaddr 10.10.11.2
```

The IP address in this example indicates the remote endpoint IP address of the FCIP tunnel.

```
MDS1(config-if)# no shut
MDS1(config-if)# end
```

The output from issuing the **show run** command displays the default values in the following example.

```
MDS1# show run

Building Configuration ...
 fcip profile 28
 ip address 10.10.10.2
 port 3225
 tcp keepalive-timeout 60
 tcp max-retransmissions 4
 tcp pmtu-enable reset-timeout 3600
 tcp initial-retransmit-time 100
 tcp window-size 64

vsan database
vsan 2 name grumpy_02

interface fcip28
no shutdown
use-profile 28
peer-info ipaddr 10.10.11.2

ip route 10.10.11.0 255.255.255.0 10.10.10.1
```

The static route must be set for FCIP tunnels. This route could also be **ip route 10.10.11.0 255.255.255.0 int gig 2/8**.

```
ips heartbeat
ips hapreset
ips boot
 interface GigabitEthernet2/8
ip address 10.10.10.2 255.255.255.0
(This is the IP address used by the FCIP profile.)
no shutdown
```

The following example shows the configuration of MDS2 with debug mode activated.

```
MDS2(config)# fcip profile 28
Mar 10 21:41:04 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32222)
Mar 10 21:41:04 ips: Create Entity 28
Mar 10 21:41:04 ips: entity28: add to config pss
```

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

```

MDS2(config-profile)# ip address 10.10.11.2
Mar 10 21:41:15 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32258)
Mar 10 21:41:15 ips: entity28: IP address changed to 10.10.11.2
Mar 10 21:41:15 ips: entity28: IP 10.10.11.2 configured for interface GigabitEthernet2/8
Mar 10 21:41:15 ips: entity28: Apply the entity config and save to config pss
Mar 10 21:41:15 ips: entity28: add to config pss

MDS2(config-profile)# exit

MDS2(config)# interface fcip 28
Mar 10 21:41:46 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32358)
Mar 10 21:41:46 ips: Verified FCIP28 Create:0
Mar 10 21:41:46 ips: FCIP28: Verified Create:0
Mar 10 21:41:46 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32360)
Mar 10 21:41:46 ips: FCIP28: Creating FCIP tunnel
Mar 10 21:41:46 ips: FCIP28: add to admin pss
Mar 10 21:41:46 ips: FCIP28: add to run-time pss
Mar 10 21:41:46 ips: FCIP28: log: 0 phy: 0 state: 0 syslog: 0

MDS2(config-if)# use-profile 28
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32480)
Mar 10 21:42:23 ips: FCIP28: Process tunnel configuration event
Mar 10 21:42:23 ips: FCIP28: Change Entity-id from 0 to 28
Mar 10 21:42:23 ips: FCIP: Optimal IF lookup for GigabitEthernet2/8 is GigabitEthernet2/8
Mar 10 21:42:23 ips: FCIP28: bind with GigabitEthernet2/8 (phy GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: Queueing bind tunnel to src if event to tunnel FSM resource:
0
Mar 10 21:42:23 ips: Locked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32480)
Mar 10 21:42:23 ips: FCIP28: Send bind for GigabitEthernet2/8 to PM (phy
GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 0 syslog: 0
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_IPS_CFG_FCIP_IF(mts opc 1905, msg id 7304)
Mar 10 21:42:23 ips: Hndlr MTS_OPC_IPS_CFG_FCIP_IF (mts_opc 1905 msg_id 7304)
Mar 10 21:42:23 ips: FCIP28: Got a tunnel param pull request from LC
Mar 10 21:42:23 ips: Added to pending queue event-id [29] event-cat [2]
Mar 10 21:42:23 ips: FCIP28: Queueing Process a Pull Request event to Pending queue
resource: 0
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_PM_FCIP_BIND(mts opc 335, msg id 32495)
Mar 10 21:42:23 ips: Hndlr MTS_OPC_PM_FCIP_BIND (mts_opc 335 msg_id 32495)
Mar 10 21:42:23 ips: FCIP28: Success received from PM for bind to GigabitEthernet2/8 (phy
GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: Bind-resp event processing bind...
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:42:23 ips: FCIP28: Last reference...
Mar 10 21:42:23 ips: FCIP28: Update the tunnel param and save to PSS
Mar 10 21:42:23 ips: FCIP28: add to admin pss
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:42:23 ips: Unlocked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32480)
Mar 10 21:42:23 ips: Dequeued pending queue msg event_id [29] cat [2]
Mar 10 21:42:23 ips: (ips_demux) Mts Opcode is 1905, id is 7304
Mar 10 21:42:23 ips: FCIP28: Processing Pull Config Request
Mar 10 21:42:23 ips: FCIP28: Bound to entity 28 port: 3225 ip: 10.10.11.2

```

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

```
MDS2(config-if)# peer-info ipaddr 10.10.10.2
Mar 10 21:43:01 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32616)
Mar 10 21:43:01 ips: FCIP28: Process tunnel configuration event
Mar 10 21:43:01 ips: FCIP28: Change Peer IP from 0.0.0.0 to 10.10.10.2 and port from 3225
to 3225
Mar 10 21:43:01 ips: FCIP28: Queueing Set tunnel param event to tunnel FSM resource: 0
Mar 10 21:43:01 ips: Locked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32616)
Mar 10 21:43:01 ips: FCIP28: Send tunnel params to LC to DPP: 7
Mar 10 21:43:01 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_SET_LC_TUNNEL_PARAM(mts opc 1897,
msg id 7358)
Mar 10 21:43:01 ips: Hndlr MTS_OPC_IPS_FCIP_SET_LC_TUNNEL_PARAM (mts_opc 1897 msg_id 7358)
Mar 10 21:43:01 ips: In handler : Received resp code: 0
Mar 10 21:43:01 ips: FCIP28: Received the tunnel params from LC
Mar 10 21:43:01 ips: FCIP28: Update the tunnel param and save to PSS
Mar 10 21:43:01 ips: FCIP28: add to admin pss
Mar 10 21:43:01 ips: FCIP28: add to run-time pss
Mar 10 21:43:01 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:43:01 ips: Unlocked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32616)

MDS2(config-if)#
MDS2(config-if)# no shut
MDS2(config-if)# Mar 10 21:43:32 ips: Dequeued mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc 3114, msg id 32737)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32737)
Mar 10 21:43:32 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 32778)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32778)
Mar 10 21:43:32 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 32783)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32783)
```

The following example shows the debug output from the supervisor of the FCIP tunnel.

```
MDS2(config)# int fcip 28
MDS2(config-if)# no shut
MDS2(config-if)# Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call
- found data in FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(0)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(0),
empty
Mar 10 22:59:46 ips: Starting a new round
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47540)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47540)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47540) dropped
Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call - found data in
FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(6)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(3),
empty
```

*Send comments to mdsfeedback-doc@cisco.com.*

```
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47589)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47589)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null_fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47589) dropped
Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call - found data in
FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(4)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(2),
empty
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47602)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47602)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null_fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47602) dropped
```

The following example shows the debug output from the IPS module of the FCIP tunnel.

MDS2# **attach module 2**

module-2# **debug ips fcip fsm port 8**

(This is the Gigabit Ethernet port 2/8.)

```
Mar 13 19:18:19 port8: 2700:FCIP28: Received new TCP connection from peer:
10.10.10.2:65455
Mar 13 19:18:19 port8: 2701:FCIP: (fcip_de_create): DE = 0xdc02ca40
Mar 13 19:18:19 port8: 2702:FCIP28: Create a DE 0xdc02ca40 for this tunnel
Mar 13 19:18:19 port8: 2703:FCIP28: Bind the DE 0xdc02ca40 [1] to tunnel LEP 0x801ebac0
Mar 13 19:18:19 port8: 2704:FCIP28: Bind DE 1 to TCP-hdl 0xdc489800
Mar 13 19:18:19 port8: 2705:FCIP28: Bind DE 1 to eport 0x801eaaa0
Mar 13 19:18:19 port8: 2706:FCIP28: bind de 1 in eport 0x801eaaa0, hash = 1 num-conn: 2
Mar 13 19:18:19 port8: 2707:FCIP28: Received new TCP connection from peer:
10.10.10.2:65453
Mar 13 19:18:19 port8: 2708:FCIP: (fcip_de_create): DE = 0xdc02cb40
Mar 13 19:18:19 port8: 2709:FCIP28: Create a DE 0xdc02cb40 for this tunnel
Mar 13 19:18:19 port8: 2710:FCIP28: Bind the DE 0xdc02cb40 [2] to tunnel LEP 0x801ebac0
Mar 13 19:18:19 port8: 2711:FCIP28: Bind DE 2 to TCP-hdl 0xdc488800
Mar 13 19:18:19 port8: 2712:FCIP28: Bind DE 2 to eport 0x801eaaa0
Mar 13 19:18:19 port8: 2713:FCIP28: bind de 2 in eport 0x801eaaa0, hash = 2 num-conn: 2
Mar 13 19:18:19 port8: 2714:FCIP28: Send LINK UP to SUP
Mar 13 19:18:20 port8: 2715:FCIP28: *** Received eisl frame in E mode
Mar 13 19:18:20 port8: 2716:FCIP28: SUP-> Set trunk mode: 2
Mar 13 19:18:20 port8: 2717:FCIP28: Change the operational mode to TRUNK
Mar 13 19:18:20 port8: 2718:FCIP28: Tunnel bringup debounce timer callback, try to bring
up tunnel
Mar 13 19:18:20 port8: 2719:FCIP28: Tunnel is already in oper UP state, don't try to
bring up again...
```

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

Use the **show fcip profile** command to verify that the configuration of the profiles are correct. The IP address and TCP port are the ports to listen on, and both are adjustable in the FCIP profile. The example below displays all default values that are adjustable while configuring the FCIP profile.

```
MDS1# show fcip profile
```

```

ProfileId Ipaddr TcpPort

28 10.10.10.2 3225
```

```
MDS1# show fcip profile 28
```

```
FCIP Profile 28
```

```
Listen Port is 3225
```

```
TCP parameters
```

```
SACK is disabled
```

```
PMTU discover is enabled, reset timeout is 3600 sec
```

```
Keep alive is 60 sec
```

```
Minimum retransmission timeout is 100 ms
```

```
Maximum number of re-transmissions is 4
```

```
Advertised window size is 64 KB
```

Use the **show interface fcip** command to verify that the interface FCIP tunnel is established and that traffic is passing through.

```
MDS1# show interface fcip 28
```

```
FCIP28 is trunking
```

```
Hardware is GigabitEthernet
```

```
Port WWN is 20:5e:00:05:30:00:59:de
```

```
Peer port WWN is 20:5e:00:0b:5f:d5:9f:c0
```

```
Admin port mode is auto, trunk mode is on
```

```
Port mode is TE
```

(The FCIP tunnel will be either E (ISL or TE (EISL) passing through multiple VSANs.)

```
vsan is 1
```

```
Trunk vsans (allowed active) (1-2)
```

```
Trunk vsans (operational) (1-2)
```

```
Trunk vsans (up) (1-2)
```

```
Trunk vsans (isolated) ()
```

```
Trunk vsans (initializing) ()
```

```
Using Profile id 28 (interface GigabitEthernet2/8)
```

(This is the FCIP profile and the Gigabit Ethernet being used by the FCIP tunnel.)

```
Peer Information
```

```
Peer Internet address is 10.10.11.2 and port is 3225
```

(This is the remote end point's IP address and listening port.)

```
Special Frame is disabled
```

(The Special Frame for verification of a remote MDS is not being used.)

```
Maximum number of TCPconnections is 2
```

(The default is 2 TCP connection being used, one for class F and other for class 2 and 3.)

```
Time Stamp is disabled
```

(The timestamp can be activated under the interface FCIP.)

```
B-port mode disabled
```

```
TCP Connection Information
```

```
2 Active TCP connections
```

```
Control connection: Local 10.10.10.2:3225, Remote 10.10.11.2:65519
```

(The above is class F traffic.)

```
Data connection: Local 10.10.10.2:3225, Remote 10.10.11.2:65521
```

(The above is class 2,3 traffic.)



*Send comments to mdsfeedback-doc@cisco.com.*

```

6 Attempts for active connections, 3 close of connections
TCP Parameters
 Path MTU 1500 bytes
 Current retransmission timeout is 100 ms <<< Default, adjusted under
 Round trip time: Smoothed 10 ms, Variance: 5

```

(This is the calculated round trip time of the FCIP tunnel. Large round trip times will require increasing the TCP window size under the FCIP profile.)

```

Advertized window: Current: 64 KB, Maximum: 64 KB, Scale: 1

```

(This is the local advertised TCP window size, and the default is 64 KB.)

```

Peer receive window: Current: 64 KB, Maximum: 64 KB, Scale: 1

```

(This is the remote end point advertised TCP window size.)

```

Congestion window: Current: 2 KB

```

(This is the minimum windows size used during congestion, and is not configurable.)

```

5 minutes input rate 136 bits/sec, 17 bytes/sec, 0 frames/sec
5 minutes output rate 136 bits/sec, 17 bytes/sec, 0 frames/sec
 2288 frames input, 211504 bytes
 2288 Class F frames input, 211504 bytes
 0 Class 2/3 frames input, 0 bytes
 0 Error frames
 2288 frames output, 211520 bytes
 2288 Class F frames output, 211520 bytes
 0 Class 2/3 frames output, 0 bytes
 0 Error frames 0 reass frames

```

```

MDS1# show interface fcip 28 brief

```

```

Interface Vsan Admin Admin Status Oper Profile Port-channel
 Mode Trunk
 Mode

fcip28 1 auto on trunking TE 28 --

```

```

MDS1# show interface fcip 28 counters bri

```

```

Interface Input (rate is 5 min avg) Output (rate is 5 min avg)

Rate Total Rate Total
Mbits/s Frames Mbits/s Frames

fcip28 18 0 18 0

```

(This is the frames that averaged over 5 minutes and the total count of frames since the last **clear counters** command was issued, or since the last tunnel up.)

Verify default 2 tcp connections are established for each fcip tunnel configured, one for control traffic and one for data traffic

```

MDS1# show ips stats tcp interface gigabitethernet 2/8

```

```

TCP Statistics for port GigabitEthernet2/8

```

```

Connection Stats

```

```

6 active openings, 8 accepts

```

```

6 failed attempts, 0 reset received, 8 established

```

```

Segment stats

```

```

295930 received, 1131824 sent, 109 retransmitted

```

(Excessive retransmits indicate possible core drops and/or that the TCP window size should be adjusted.)

```

0 bad segments received, 0 reset sent

```

*Send comments to mdsfeedback-doc@cisco.com.*

```
TCP Active Connections
 Local Address Remote Address State Send-Q Recv-Q
 10.10.10.2:3225 10.10.11.2:65519 ESTABLISH 0 0
(This is used for F control traffic only.)

10.10.10.2:3225 10.10.11.2:65521 ESTABLISH 87568 0
(Send-Q increasing during read-only test.)

10.10.10.2:3225 0.0.0.0:0 LISTEN 0 0
(The TCP listen port is ready for new TCP connections.)
```

You can use the following command to verify that traffic is incrementing on Gigabit Ethernet port of the FCIP tunnel.

```
MDS1# show ips stats mac interface gigabitethernet 2/8
Ethernet MAC statistics for port GigabitEthernet2/8
 Hardware Transmit Counters
 1074898 frame 1095772436 bytes
 0 collisions, 0 late collisions, 0 excess collisions
 0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
 Hardware Receive Counters
 33488196 bytes, 298392 frames, 277 multicasts, 16423 broadcasts
 0 bad, 0 runt, 0 CRC error, 0 length error
 0 code error, 0 align error, 0 oversize error
 Software Counters
 298392 received frames, 1074898 transmit frames
 0 frames soft queued, 0 current queue, 0 max queue
 0 dropped, 0 low memory
```

Traffic statistics can be verified on the internal ASIC chip on each Gigabit Ethernet port.

```
MDS1# show ips stats flamingo interface gigabitethernet 2/8
Flamingo ASIC Statistics for port GigabitEthernet2/8
 Hardware Egress Counters
 2312 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
(Good frames and CRC error frames can be monitored.)

 Hardware Ingress Counters
(Verify good increments on the active tunnel.)

 2312 Good, 0 protocol error, 0 header checksum error
 0 FC CRC error, 0 iSCSI CRC error, 0 parity error
 Software Egress Counters
 2312 good frames, 0 bad header cksum, 0 bad FIFO SOP
 0 parity error, 0 FC CRC error, 0 timestamp expired error
 0 unregistered port index, 0 unknown internal type
 0 RDL, 0 RDL too big RDL, 0 TDL ttl_1
 3957292257 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
 Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
 Software Ingress Counters
 2312 Good frames, 0 header cksum error, 0 FC CRC error
 0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
 0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
 0 out of memory drop, 0 queue full drop
 0 RDL, 0 too big RDL drop
 Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
```

On the next few pages are screen captures taken with Ethereal, of TCP connection being established, and FCIP tunnels. Note that FCIP tunnel activation is the same as an FC EISL becoming active (such as ELP, ESC, and EFP). The following traces were captured after configuration on both MDS 9000 Family switches, and the last “no shutdown” was entered on switch MDS1. All settings are default (for example, SACK is disabled, the TCP window is set to 64K).

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

Figure 5-3 First Capture of TCP Connection

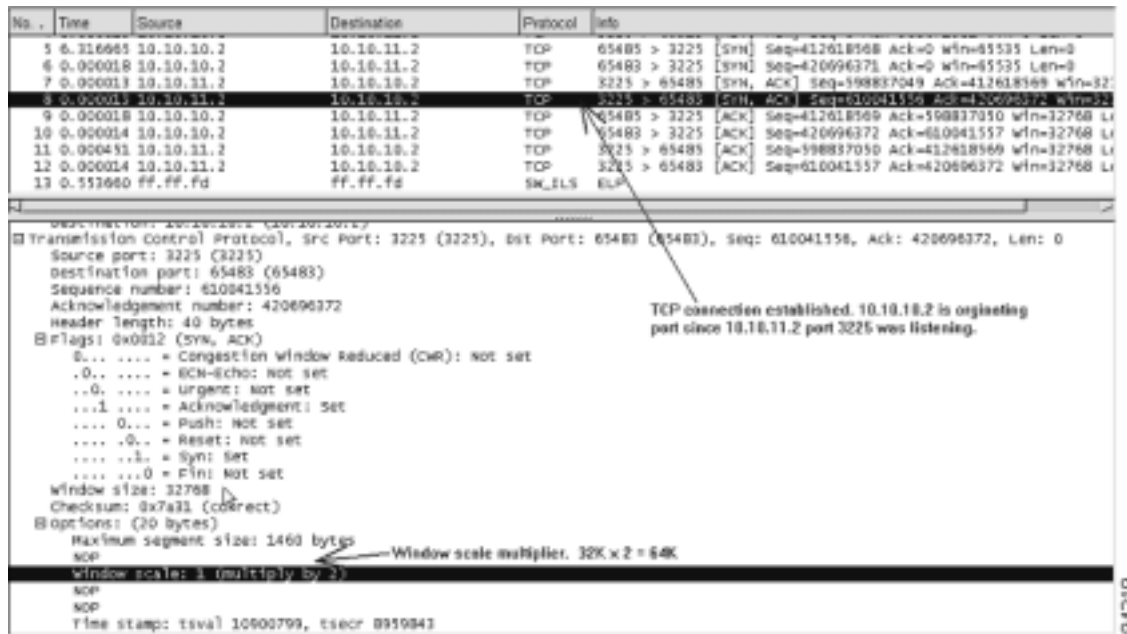
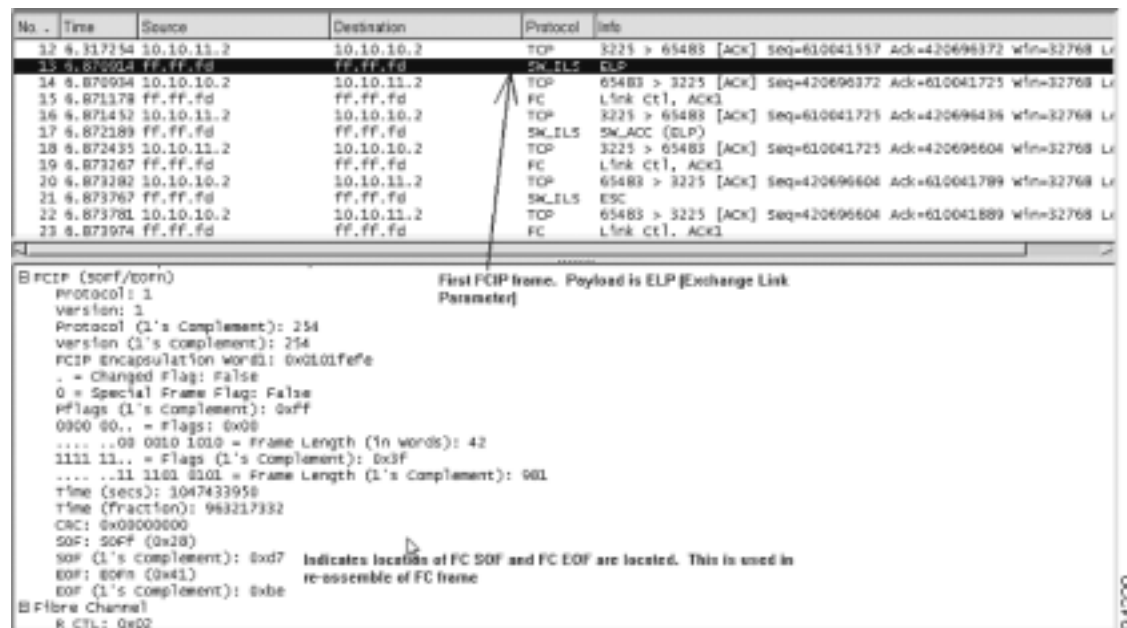


Figure 5-4 shows more of the trace, with frame 13 being the first FCIP frame. This frame carries the FC Standard ELP.

Figure 5-4 Second Capture of TCP Connection



*Send comments to mdsfeedback-doc@cisco.com.*

Figure 5-5 shows the FC portion of the EISL initialization over the FCIP tunnel.

**Figure 5-5 Third Capture of TCP Connection**

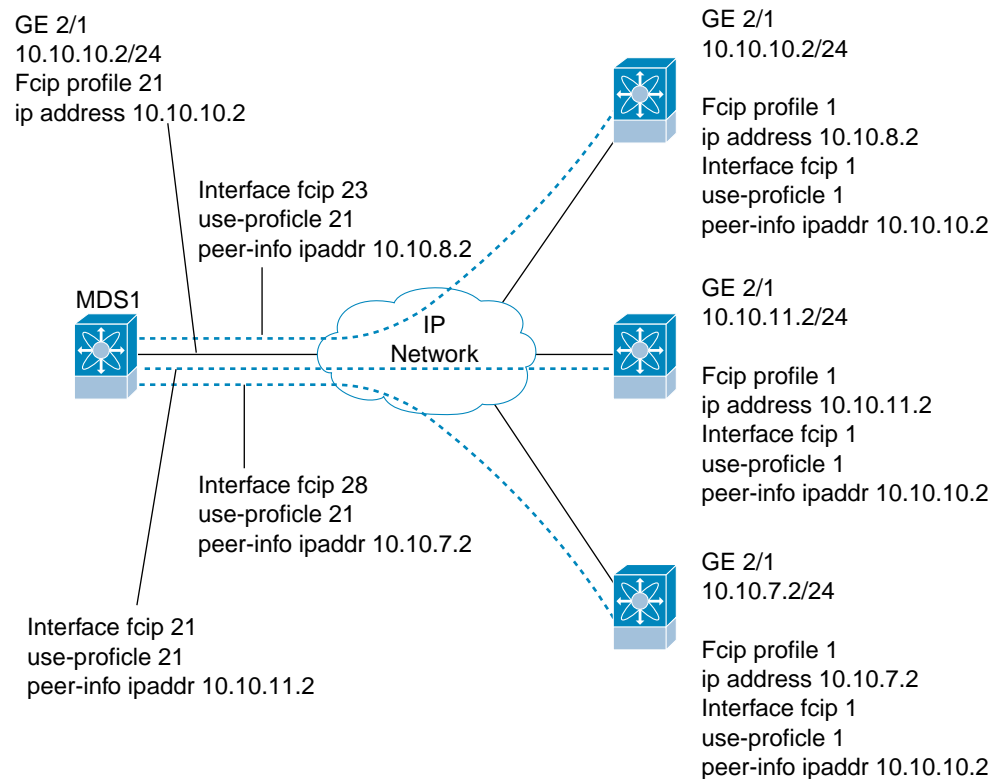
| File | Edit     | Capture    | Display     | Tools    | Help                                      |
|------|----------|------------|-------------|----------|-------------------------------------------|
| No.  | Time     | Source     | Destination | Protocol | Info                                      |
| 13   | 0.000000 | ff.ff.fd   | ff.ff.fd    | TCP      | 65483 → 3225 [ACK] Seq=420696372 Ack=6100 |
| 14   | 0.000000 | 10.10.10.2 | 10.10.11.2  | SN_ILS   | ELP                                       |
| 15   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 16   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610041725 Ack=4206 |
| 17   | 0.000000 | ff.ff.fd   | ff.ff.fd    | SN_ILS   | SW_ACC (ELP)                              |
| 18   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610041725 Ack=4206 |
| 19   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 20   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420696604 Ack=6100 |
| 21   | 0.000000 | ff.ff.fd   | ff.ff.fd    | SN_ILS   | ESC                                       |
| 22   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420696604 Ack=6100 |
| 23   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 24   | 0.000000 | ff.ff.fd   | ff.ff.fd    | SN_ILS   | SW_ACC (ESC)                              |
| 25   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610041889 Ack=4206 |
| 26   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610041889 Ack=4206 |
| 27   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 28   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420696756 Ack=6100 |
| 29   | 0.000000 | ff.ff.fd   | ff.ff.fd    | SN_ILS   | 0x71                                      |
| 30   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420696756 Ack=6100 |
| 31   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 32   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610042169 Ack=4206 |
| 33   | 0.000000 | ff.ff.fd   | ff.ff.fd    | SN_ILS   | SW_ACC (0x71)                             |
| 34   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610042169 Ack=4206 |
| 35   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 36   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420697436 Ack=6100 |
| 37   | 0.000000 | 00.00.00   | 00.00.00    | SN_ILS   | 0x71                                      |
| 38   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610042633 Ack=4206 |
| 39   | 0.000000 | ff.ff.fd   | ff.ff.fd    | FC       | Link Ct1, ACK1                            |
| 40   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420697528 Ack=6100 |
| 41   | 0.000000 | 00.00.00   | 00.00.00    | SN_ILS   | SW_ACC (0x71)                             |
| 42   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420697528 Ack=6100 |
| 43   | 0.000000 | 00.00.00   | 00.00.00    | FC       | Link Ct1, ACK1                            |
| 44   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610042789 Ack=4206 |
| 45   | 0.000000 | 00.00.00   | 00.00.00    | SN_ILS   | SWP                                       |
| 46   | 0.000000 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 → 65483 [ACK] Seq=610042789 Ack=4206 |
| 47   | 0.000000 | 00.00.00   | 00.00.00    | SN_ILS   | SWP                                       |
| 48   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420697688 Ack=6100 |
| 49   | 0.000000 | 00.00.00   | 00.00.00    | FC       | Link Ct1, ACK1                            |
| 50   | 0.000000 | 00.00.00   | 00.00.00    | FC       | Link Ct1, ACK1                            |
| 51   | 0.000000 | 10.10.10.2 | 10.10.11.2  | TCP      | 65483 → 3225 [ACK] Seq=420697760 Ack=6100 |

*Send comments to mdsfeedback-doc@cisco.com.*

## One to three FCIP tunnel creation and monitoring

Figure 5-6 shows the configuration of switch MDS1 for three tunnels from one Gigabit Ethernet port.

**Figure 5-6 MDS1 Configured for Three FCIP Tunnels**



The following example shows the configuration of switch MDS1 for three tunnels from one Gigabit Ethernet port.

```
MDS1(config)# fcip profile 21
MDS1(config-profile)# ip address 10.10.10.2
MDS1(config-profile)# exit
MDS1(config)# interface fcip 21
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.11.2
MDS1(config-if)# no shut
MDS1(config-if)# exit

MDS1(config)# ip route 10.10.11.0 255.255.255.0 10.10.10.1
MDS1(config)# ip route 10.10.11.0 255.255.255.0 int gigabitethernet 2/1
```

Now the interface FCIP is created for the second tunnel. The same FCIP profile is used for this example. A separate FCIP profile can be used for each interface FCIP if desired.

```
MDS1(config-if)#
MDS1(config-if)# int fcip 23
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.8.2
MDS1(config-if)# no shut
```

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

```
MDS1(config-if)# exit
MDS1(config)#
```

Now the FCIP interface is created for the third tunnel.

```
MDS1(config)# int fcip 28
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.7.2
MDS1(config-if)# no shut
MDS1(config-if)# end
```

## FCIP Profile Misconfiguration Examples

The following example shows an incorrect or non-existent IP address used for an FCIP profile.

```
MDS22(config)# fcip profile 21
MDS22(config-profile)# ip addr 1.1.1.1
MDS22(config-profile)# ip addr 34.34.34.34
MDS22(config-profile)# exit
MDS22(config)# exit
MDS22# show fcip profile 21
FCIP Profile 21
```

```
Internet Address is 34.34.34.34
```

(In the line above, the interface Gigabit Ethernet port is not shown. This means the IP address is not assigned a Gigabit Ethernet port.

```
Listen Port is 3225
TCP parameters
SACK is disabled
PMTU discover is enabled, reset timeout is 3600 sec
Keep alive is 60 sec
Minimum retransmission timeout is 300 ms
Maximum number of re-transmissions is 4
Advertised window size is 64 KB
```

```
MDS22# conf t
Enter configuration commands, one per line. End with CNTL/Z.
MDS22(config)# int gigabitethernet 2/5
MDS22(config-if)# ip addr 34.34.34.34 255.255.255.0
MDS22(config-if)# no shut
MDS22(config-if)# end
MDS22# show fcip profile 34
error: fcip profile not found
MDS22# show fcip profile 21
FCIP Profile 21
```

```
Internet Address is 34.34.34.34 (interface GigabitEthernet2/5)
```

(In the line above, the Gigabit Ethernet port is now shown and the FCIP profile is bound to a physical port.)

```
Listen Port is 3225
TCP parameters
SACK is disabled
PMTU discover is enabled, reset timeout is 3600 sec
Keep alive is 60 sec
Minimum retransmission timeout is 300 ms
Maximum number of re-transmissions is 4
Advertised window size is 64 KB
```

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

The following example shows a configuration error when using multiple FCIP profiles on one physical Gigabit Ethernet port.

```
MDS2(config)# fcip profile 21
MDS2(config-profile)# ip address 10.10.11.2
error: fcip another profile exists with same port & ip
(Multiple FCIP profiles can be used on one physical Gigabit Ethernet port, but each profile must have a
different listening port.)
```

```
MDS2(config-profile)# port ?
<1-65535> Profile TCP Port
```

```
MDS2(config-profile)# port 32
(Change the TCP listening port on the profile. The default is 3225.)
```

```
MDS2(config-profile)# ip address 10.10.11.2
(The IP address for the Gigabit Ethernet port 2/1 is now accepted, and two FCIP profiles are using the
same Gigabit Ethernet port.)
```

```
MDS2(config-profile)# end
MDS2# show fcip profile 21
FCIP Profile 21
 Internet Address is 10.10.11.2 (interface GigabitEthernet2/1)
 Listen Port is 32
```

(This is a new TCP listen port.)

```
TCP parameters
 SACK is disabled
 PMTU discover is enabled, reset timeout is 3600 sec
 Keep alive is 60 sec
 Minimum retransmission timeout is 300 ms
 Maximum number of re-transmissions is 4
 Advertised window size is 64 KB
```

```
MDS2# show fcip profile 28
FCIP Profile 28
 Internet Address is 10.10.11.2 (interface GigabitEthernet2/1)
 Listen Port is 3225
```

(This is the default listen port.)

```
TCP parameters
 SACK is disabled
 PMTU discover is enabled, reset timeout is 3600 sec
 Keep alive is 60 sec
 Minimum retransmission timeout is 300 ms
 Maximum number of re-transmissions is 4
 Advertised window size is 64 KB
```

The following example shows a configuration error when bringing a tunnel up on the selected port. This could be either an FCIP profile issue or an interface FCIP issue. Both sides must be configured correctly.

```
MDS2(config)# fcip profile 21
MDS2(config-profile)# port 13
(Change the TCP listen port on switch MDS2.)
```

```
MDS2(config-profile)# end
MDS2(config)# int fcip 21
MDS2(config-if)# passive-mode
(Put interface FCIP 21 in passive mode to guarantee MDS1 initiates a TCP connection.)
```

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

```

module-2# debug ips fcip fsm port 1
module-2# Mar 14 23:08:02 port1: 863:FCIP21: SUP-> Set Port mode 1
Mar 14 23:08:02 port1: 864:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 23:08:02 port1: 865:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 23:08:02 port1: 866:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 23:08:02 port1: 867:FCIP21: Try to Bring UP the Tunnel
Mar 14 23:08:02 port1: 868:FCIP21: Start TCP listener with peer: 10.10.10.2:13
(This debug output from switch MDS2 shows that the FCIP tunnel will not come up because switch
MDS2 is listening on port 13, and switch MDS1 is trying to establish the connection on the default port
3225.)

Mar 14 23:08:02 port1: 869:FCIP: Create a new listener object for 10.10.11.2:13
Mar 14 23:08:02 port1: 870:FCIP: Create FCIP Listener with local info: 10.10.11.2:13

```

```

MDS1(config)# int fcip 21
MDS1(config-if)# peer-info ip 10.10.11.2 port 13
(The remote end interface FCIP must be configured to establish a TCP connection on a port that is being
used as TCP listen port.)

```

```

MDS1(config-if)# end
MDS1# show int fcip 21
fcip21 is trunking
(The FCIP tunnel is now up.)

```

```

Hardware is GigabitEthernet
Port WWN is 20:42:00:05:30:00:59:de
Peer port WWN is 20:42:00:0b:5f:d5:9f:c0
Admin port mode is auto, trunk mode is on
Port mode is TE
vsan is 1
Trunk vsans (allowed active) (1-2)
Trunk vsans (operational) (1-2)
Trunk vsans (up) ()
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1-2)
Using Profile id 21 (interface GigabitEthernet2/1)
Peer Information
Peer Internet address is 10.10.11.2 and port is 13
Special Frame is disabled
Maximum number of TCP connections is 2
Time Stamp is disabled
B-port mode disabled
TCP Connection Information
2 Active TCP connections
Control connection: Local 10.10.10.2:65188, Remote 10.10.11.2:13
(The port is 13 as configured.)

Data connection: Local 10.10.10.2:65190, Remote 10.10.11.2:13
174 Attempts for active connections, 5 close of connections

```



*Send comments to mdsfeedback-doc@cisco.com.*

```
MDS2# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
 Connection Stats
 44 active openings, 2 accepts
 26 failed attempts, 0 reset received, 20 established
 Segment stats
 2515 received, 2342 sent, 0 retransmitted
 0 bad segments received, 0 reset sent

 TCP Active Connections
 Local Address Remote Address State Send-Q Recv-Q
 10.10.11.2:13 10.10.10.2:65188 ESTABLISH 0 0
(The port is 13 as configured.)
 10.10.11.2:13 10.10.10.2:65190 ESTABLISH 0 0
(The port is 13 as configured.)
 10.10.11.2:13 0.0.0.0:0 LISTEN 0 0
 0.0.0.0:3260 0.0.0.0:0 LISTEN 0 0
```

## Interface FCIP Misconfiguration Examples

The following example shows the “peer-info” IP address of the remote end-point is missing. The debug output is from the IPS module.

```
Module-2# debug ips fcip fsm port 1
module-2# Mar 14 21:37:05 port1: 38:FCIP21: SUP-> Set Port mode 1
Mar 14 21:37:05 port1: 39:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 21:37:05 port1: 40:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 21:37:05 port1: 41:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 21:37:05 port1: 42:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:37:05 port1: 43:FCIP21: Bring up tunnel Failed, peer-ip not set
(The peer IP address is not set.)
```

```
MDS2# show int fcip 21
fcip21 is down (Link failure or not-connected)
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Admin port mode is auto, trunk mode is on
 vsan is 1
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
```

(This line shows the Peer Information as empty. The line should read “Peer Internet address is 10.10.10.2 and port is 3225.”)

```
 Special Frame is disabled
 Maximum number of TCP connections is 2
 Time Stamp is disabled
 B-port mode disabled
 TCP Connection Information
 0 Attempts for active connections, 0 close of connections
 5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 0 frames input, 0 bytes
 0 Class F frames input, 0 bytes
 0 Class 2/3 frames input, 0 bytes
 0 Error frames
 0 frames output, 0 bytes
 0 Class F frames output, 0 bytes
 0 Class 2/3 frames output, 0 bytes
 0 Error frames 0 reass frames
```

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

The following example shows the interface FCIP is administratively shut down. The debug output is from the IPS module.

```
Module-2# debug ips fcip fsm port 1
module-2# Mar 14 21:32:27 port1: 1:FCIP21: Create tunnel with ifindex: a000014
Mar 14 21:32:27 port1: 2:FCIP21: Get the peer info from the SUP-IPS-MGR
Mar 14 21:32:27 port1: 3:FCIP21: SUP-> Disable tunnel: already in disable state
Mar 14 21:32:27 port1: 4:FCIP21: SUP-> Set Port mode 1
Mar 14 21:32:27 port1: 5:FCIP21: SUP-> Set port index: 21
Mar 14 21:32:27 port1: 6:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 7:FCIP21: Tunnel in admin down state
(The tunnel needs no shut down on the interface FCIP.)

Mar 14 21:32:27 port1: 8:FCIP21: SUP-> Set port VSAN: 1
Mar 14 21:32:27 port1: 9:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 10:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 11:FCIP21: SUP-> Set port WWN: 0x2042000b5fd59fc0
Mar 14 21:32:27 port1: 12:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 13:FCIP21: Tunnel in admin down state
(The tunnel needs no shut down on the interface FCIP.)

Mar 14 21:32:27 port1: 14:FCIP21: SUP-> Set trunk mode: 1
Mar 14 21:32:27 port1: 15:FCIP21: SUP-> Set source IF: 2080000
Mar 14 21:32:27 port1: 16:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 17:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 18:FCIP21: SUP-> Switch WWN: 0x2000000b5fd59fc0
Mar 14 21:32:27 port1: 19:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 20:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 21:FCIP21: SUP-> Response to SB's pull all tunnel info
Mar 14 21:32:27 port1: 22:FCIP21: SUP-> Set peer port: 3225 current port: 3225
Mar 14 21:32:27 port1: 23:FCIP21: peer port has same value, do nothing
Mar 14 21:32:27 port1: 24:FCIP21: Set number of tcp connection 2
Mar 14 21:32:27 port1: 25:FCIP21: SUP-> Set Local listen IP: 10.10.11.2 current ip
0.0.0.0
Mar 14 21:32:27 port1: 26:FCIP21: SUP-> Set Local listen Port: 3225 current port 3225
Mar 14 21:32:27 port1: 27:FCIP21: SUP-> Enable PMTU Discovery, timeout 3600
Mar 14 21:32:27 port1: 28:FCIP21: SUP-> Set round-trip time to 300 ms. Current value 100
ms
Mar 14 21:32:27 port1: 29:FCIP21: SUP-> Set keep-alive time to 60 sec. current value 60
sec
```

```
MDS2# show int fcip 21
fcip21 is down (Administratively down)
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Admin port mode is auto, trunk mode is on
 vsan is 1
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.10.2 and port is 3225
 Special Frame is disabled
 Maximum number of TCP connections is 2
```

Local MDS trying to connect to remote end point on port 13 and remote end point set to default listen port 3225

```
MDS2# show int fcip 21
fcip21 is down (Link failure or not-connected)
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Admin port mode is auto, trunk mode is on
 vsan is 1
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.10.2 and port is 13
```

*Send comments to mdsfeedback-doc@cisco.com.*

```
MDS1# show fcip profile 21
FCIP Profile 21
 Internet Address is 10.10.10.2 (interface GigabitEthernet2/1)
 Listen Port is 3225
 TCP parameters
 SACK is disabled
 PMTU discover is enabled, reset timeout is 3600 sec
 Keep alive is 60 sec
 Minimum retransmission timeout is 300 ms
 Maximum number of re-transmissions is 4
 Advertised window size is 64 KB
```

The following debug output is from switch MDS2.

```
Mar 14 23:26:07 port1: 1340:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Mar 14 23:26:07 port1: 1341:FCIP: Create a new listener object for 10.10.11.2:3225
Mar 14 23:26:07 port1: 1342:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Mar 14 23:26:07 port1: 1343:FCIP21: Create a DE 0xd802d140 for this tunnel
Mar 14 23:26:07 port1: 1344:FCIP21: Bind the DE 0xd802d140 [1] to tunnel LEP 0x80111570
Mar 14 23:26:07 port1: 1345:FCIP21: Start the active connection [1] to 10.10.10.2:13
Mar 14 23:26:07 port1: 1346:FCIP21: Create a DE 0xd802cdc0 for this tunnel
Mar 14 23:26:07 port1: 1347:FCIP21: Bind the DE 0xd802cdc0 [2] to tunnel LEP 0x80111570
Mar 14 23:26:07 port1: 1348:FCIP21: Start the active connection [2] to 10.10.10.2:13
(The switch is attempting to create a TCP connection on port 13. The creation port must match the TCP
listen port on the remote end point.)
```

```
Mar 14 23:26:07 port1: 1349:FCIP21: Active Connect creation FAILED [1]
Mar 14 23:26:07 port1: 1350:FCIP21: Delete the DE [1]0xd802d140
Mar 14 23:26:07 port1: 1351:FCIP21: Delete the DE object [1] 0xd802d140
Mar 14 23:26:07 port1: 1352:FCIP21: Try 7 to bring up the tunnel
Mar 14 23:26:07 port1: 1353:FCIP21: Start the bringup tunnel timer, timeout: 64000
Mar 14 23:26:07 port1: 1354:FCIP21: Active Connect creation FAILED [2]
Mar 14 23:26:07 port1: 1355:FCIP21: Delete the DE [2]0xd802cdc0
Mar 14 23:26:07 port1: 1356:FCIP21: Set lep operation state to DOWN
Mar 14 23:26:07 port1: 1357:FCIP21: Delete the DE object [2] 0xd802cdc0
Mar 14 23:26:07 port1: 1358:FCIP21: Try 8 to bring up the tunnel
Mar 14 23:26:07 port1: 1359:FCIP21: Start the bringup tunnel timer, timeout: 128000
```

```
MDS2(config-if)# peer-info ipaddr 10.10.10.2 port 3225
(This changes the start active connection port to match the default port 3225.)
```

Or you can use this command:

```
MDS2(config-if)# no peer-info ipaddr 10.10.10.2 port 13
(Removing port 13 will also set it to the default of 3225.)
```

```
MDS2# show int fcip 21
fcip21 is trunking
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Peer port WWN is 20:42:00:05:30:00:59:de
 Admin port mode is auto, trunk mode is on
 Port mode is TE
 vsan is 1
 Trunk vsans (allowed active) (1-2)
 Trunk vsans (operational) (1-2)
 Trunk vsans (up) (1-2)
 Trunk vsans (isolated) ()
 Trunk vsans (initializing) ()
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.10.2 and port is 3225
```

*Send comments to mdsfeedback-doc@cisco.com.*

```

Special Frame is disabled
Maximum number of TCP connections is 2
Time Stamp is disabled
B-port mode disabled
TCP Connection Information
 2 Active TCP connections
 Control connection: Local 10.10.11.2:65330, Remote 10.10.10.2:3225
 Data connection: Local 10.10.11.2:65332, Remote 10.10.10.2:3225

```

In the following example, passive mode is set on both sides of the FCIP tunnel.

```

module-2# Mar 14 23:49:06 port1: 1870:FCIP21: SUP-> Set Port mode 1
Mar 14 23:49:06 port1: 1871:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 23:49:06 port1: 1872:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 23:49:06 port1: 1873:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 23:49:06 port1: 1874:FCIP21: Try to Bring UP the Tunnel
Mar 14 23:49:06 port1: 1875:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Mar 14 23:49:06 port1: 1876:FCIP: Create a new listener object for 10.10.11.2:3225
Mar 14 23:49:06 port1: 1877:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Mar 14 23:49:06 port1: 1878:FCIP21: Passive mode set, don't initiate TCP connection
(A TCP connection will not be established when passive mode is set. The Gigabit Ethernet port will only listen.)

```

```

MDS2# show int fcip 21
fcip21 is down (Link failure or not-connected)
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Admin port mode is auto, trunk mode is on
 vsan is 1
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.10.2 and port is 3225
 Passive mode is enabled

```

(Passive mode is set, so a TCP connection will not be established.)

```

Special Frame is disabled
MDS1# show int fcip 21
fcip21 is down (Link failure or not-connected)
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:05:30:00:59:de
 Admin port mode is auto, trunk mode is on
 vsan is 1
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.11.2 and port is 3225
 Passive mode is enabled

```

(Both sides are set to passive mode. You must change one or both sides to **no passive-mode** under the interface FCIP.)

```

Special Frame is disabled
MDS2(config)# int fcip 21
MDS2(config-if)# no passive-mode
(Change one or both sides to no passive-mode.)

```

```

MDS2# show int fcip 21
fcip21 is trunking

```

*Send comments to mdsfeedback-doc@cisco.com.*

The following example shows a time stamp acceptable difference failure, or no NTP server connected to synchronize clocks. When using time stamps, the MDS switch must be a synchronized clock. NTP is configurable on the MDS 9000 switch.

```
MDS2(config)# int fcip 21
MDS2(config-if)# time-stamp

module-2# debug ips fcip fsm port 1
Mar 15 00:01:35 port1: 3248:FCIP21: IPS-> Enable timestamp acceptable difference 1000
(Timestamp is enabled under the interface FCIP. The default acceptable difference is 1000.)

Mar 15 00:01:35 port1: 3249:FCIP21: IPS-> acc diff in sec: 0x1 frac: 0x0
Mar 15 00:01:35 port1: 3250:FCIP21: Sending response code: 0
Mar 15 00:01:48 port1: 3251:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
(The timestamp difference failed the acceptable difference.)

Mar 15 00:01:48 port1: 3252:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3253:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
<<< cut >>>
Mar 15 00:01:48 port1: 3290:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3291:FCIP21: (fcip_de_rcv): Previous partial packet -
Concatenating
Mar 15 00:01:48 port1: 3292:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3293:FCIP21: FCIP frame len 0x300 is not within correct range <<<
?? >>>
Mar 15 00:01:48 port1: 3294:FCIP21: Delete the DE [2]0xd802d680
Mar 15 00:01:48 port1: 3295:FCIP21: replace the eport entry at index: 1
Mar 15 00:01:48 port1: 3296:FCIP21: DE [-670902656] 0x00000002 terminate tcp connection
0xd8072800
(The TCP connection is disconnected because the timestamp difference is too large.)

Mar 15 00:01:48 port1: 3297:FCIP21: Delete the DE object [2] 0xd802d680
Mar 15 00:01:48 port1: 3298:FCIP21: Delete the DE [1]0xd802cf00
Mar 15 00:01:48 port1: 3299:FCIP21: Unregister from flamingo port_index: 0x21
Mar 15 00:01:48 port1: 3300:FCIP21: Send Link down to SUP
Mar 15 00:01:48 port1: 3301:FCIP21: Start the bringup tunnel timer, timeout: 18470
Mar 15 00:01:48 port1: 3302:FCIP21: replace the eport entry at index: 0
Mar 15 00:01:48 port1: 3303:FCIP21: Set lep operation state to DOWN
Mar 15 00:01:48 port1: 3304:FCIP21: DE [-670904576] 0x00000001 terminate tcp connection
0xd8072c00
Mar 15 00:01:48 port1: 3305:FCIP21: Delete the DE object [1] 0xd802cf00
Mar 15 00:01:50 port1: 3306:FCIP21: Received new TCP connection from peer:
10.10.10.2:65066
(The TCP connection begins trying to re-establish the connection.)

Mar 15 00:01:50 port1: 3307:FCIP21: Tunnel is not ADMIN UP state, reject new TCP
connection from 10.10.10.2:65066
Mar 15 00:01:50 port1: 3308:FCIP21: Received new TCP connection from peer:
10.10.10.2:65064
Mar 15 00:01:50 port1: 3309:FCIP21: Tunnel is not ADMIN UP state, reject new TCP
connection from 10.10.10.2:65064
Mar 15 00:01:56 port1: 3310:FCIP21: SUP-> Set Port mode 1
Mar 15 00:01:56 port1: 3311:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 15 00:01:56 port1: 3312:FCIP21: SUP-> Set trunk mode: 1
Mar 15 00:01:56 port1: 3313:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 15 00:01:56 port1: 3314:FCIP21: Try to Bring UP the Tunnel
Mar 15 00:01:56 port1: 3315:FCIP21: tunnel bring-up debounce timer set, wait for timer to
pop
(Connect the NTP server or synchronized clocks, or increase the acceptable difference.)
```

*Send comments to mdsfeedback-doc@cisco.com.*

```

module-2# debug ips fcip fsm port 1
module-2#
Jan 14 14:22:08 port1: 854886:FCIP21: IPS-> Enable timestamp acceptable difference 2000
Jan 14 14:22:08 port1: 854887:FCIP21: IPS-> acc diff in sec: 0x2 frac: 0x0
(The timestamp acceptable difference passes and the tunnel continues to be brought up.)

module-2#
module-2# Jan 14 14:22:39 port1: 854932:FCIP21: Received new TCP connection from peer:
10.10.10.2:64172
Jan 14 14:22:39 port1: 854933:FCIP21: Create a DE 0xd802d5c0 for this tunnel
Jan 14 14:22:39 port1: 854934:FCIP21: Bind the DE 0xd802d5c0 [1] to tunnel LEP 0x80111570
Jan 14 14:22:39 port1: 854935:FCIP21: Bind DE 1 to TCP-hdl 0xd8071000
Jan 14 14:22:39 port1: 854936:FCIP21: Bind DE 1 to eport 0x80110550
Jan 14 14:22:39 port1: 854937:FCIP21: bind de 1 in eport 0x80110550, hash = 1 num-conn: 2
Jan 14 14:22:39 port1: 854938:FCIP21: Received new TCP connection from peer: 10
.10.10.2:64170
Jan 14 14:22:39 port1: 854939:FCIP21: Create a DE 0xd802c900 for this tunnel
Jan 14 14:22:39 port1: 854940:FCIP21: Bind the DE 0xd802c900 [2] to tunnel LEP
0x80111570
Jan 14 14:22:39 port1: 854941:FCIP21: Bind DE 2 to TCP-hdl 0xd8070000
Jan 14 14:22:39 port1: 854942:FCIP21: Bind DE 2 to eport 0x80110550
Jan 14 14:22:39 port1: 854943:FCIP21: bind de 2 in eport 0x80110550, hash = 2 n
um-conn: 2
Jan 14 14:22:39 port1: 854944:FCIP21: Send LINK UP to SUP
Jan 14 14:22:39 port1: 854945:FCIP21: *** Received eisl frame in E mode
Jan 14 14:22:39 port1: 854946:FCIP21: SUP-> Set trunk mode: 2
Jan 14 14:22:39 port1: 854947:FCIP21: Change the operational mode to TRUNK

MDS2# show int fcip 21
fcip21 is trunking
 Hardware is GigabitEthernet
 Port WWN is 20:42:00:0b:5f:d5:9f:c0
 Peer port WWN is 20:42:00:05:30:00:59:de
 Admin port mode is auto, trunk mode is on
 Port mode is TE
 vsan is 1
 Trunk vsans (allowed active) (1-2)
 Trunk vsans (operational) (1-2)
 Trunk vsans (up) (1-2)
 Trunk vsans (isolated) ()
 Trunk vsans (initializing) ()
 Using Profile id 21 (interface GigabitEthernet2/1)
 Peer Information
 Peer Internet address is 10.10.10.2 and port is 3225
 Special Frame is disabled
 Maximum number of TCP connections is 2
 Time Stamp is enabled, acceptable time difference 2000 ms
 B-port mode disabled
 TCP Connection Information

```

*Send comments to mdsfeedback-doc@cisco.com.*

Figure 5-7 shows a trace of timestamp difference failure.

**Figure 5-7 Trace of Timestamp Difference Failure**

| No. | Time     | Source            | Destination       | Protocol | Info                                                          |
|-----|----------|-------------------|-------------------|----------|---------------------------------------------------------------|
| 16  | 2.102222 | ff.ff.fd          | ff.ff.fd          | SW_ILS   | ELP                                                           |
| 17  | 2.102596 | 10.10.11.2        | 10.10.10.2        | TCP      | 3225 > 64136 [ACK] Seq=863425371 Ack=833197805 wln=32768      |
| 18  | 2.251899 | 00:03:fe:6f:67:fe | 01:00:0c:cc:cc:cc | CDP      | Cisco Discovery Protocol                                      |
| 19  | 4.152642 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [FIN, ACK] Seq=833197805 Ack=863425371 wln=32768 |
| 20  | 4.152694 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [FIN, ACK] Seq=833197805 Ack=863425371 wln=32768 |
| 21  | 4.152677 | 10.10.11.2        | 10.10.10.2        | TCP      | 3225 > 64136 [FIN, ACK] Seq=863425371 Ack=833197806 wln=32768 |
| 22  | 4.152690 | 10.10.11.2        | 10.10.10.2        | TCP      | 3225 > 64136 [FIN, ACK] Seq=863425371 Ack=833197806 wln=32768 |
| 23  | 4.152709 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [ACK] Seq=833197806 Ack=863425372 wln=32768      |
| 24  | 4.152723 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [ACK] Seq=833197806 Ack=863425372 wln=32768      |
| 25  | 4.162109 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [SYN] Seq=950846288 Ack=0 wln=65535 Len=0        |
| 26  | 4.162124 | 10.10.10.2        | 10.10.11.2        | TCP      | 64136 > 3225 [SYN] Seq=950846288 Ack=0 wln=65535 Len=0        |
| 27  | 4.162137 | 10.10.11.2        | 10.10.10.2        | TCP      | 3225 > 64136 [SYN, ACK] Seq=950846288 Ack=863425371 wln=32768 |
| 28  | 4.162150 | 10.10.11.2        | 10.10.10.2        | TCP      | 3225 > 64136 [SYN, ACK] Seq=950846288 Ack=863425371 wln=32768 |

|                                                                                                                                                                                                                                                                                                                                                                                                                   |  |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| Transmission Control Protocol, Src Port: 64136 (64136), Dst Port: 3225 (3225), Seq: 833197805, Ack: 863425371, Len: 0                                                                                                                                                                                                                                                                                             |  |
| Source port: 64136 (64136)                                                                                                                                                                                                                                                                                                                                                                                        |  |
| Destination port: 3225 (3225)                                                                                                                                                                                                                                                                                                                                                                                     |  |
| Sequence number: 833197805                                                                                                                                                                                                                                                                                                                                                                                        |  |
| Acknowledgement number: 863425371                                                                                                                                                                                                                                                                                                                                                                                 |  |
| Header length: 32 bytes                                                                                                                                                                                                                                                                                                                                                                                           |  |
| Flags: 0x0011 (FIN, ACK)                                                                                                                                                                                                                                                                                                                                                                                          |  |
| 0... .. = Congestion window reduced (CWR): Not set<br>.0... .. = ECN-Echo: Not set<br>..0... .. = Urgent: Not set<br>....1... .. = Acknowledgment: set<br>....0... .. = Push: Not set<br>....0... .. = Reset: Not set<br>....0... .. = Syn: Not set<br>....1... .. = Fin: set<br>window size: 32768<br>Checksum: 0x51f1 (correct)<br>Options: (12 bytes)<br>NOP<br>NOP<br>Time stamp: tsval 8586898, tsecr 433677 |  |

MDS1 receives MDS2 ELP. Timestamp is enabled and time difference is beyond acceptable difference

TCP FIN, ACK is sent to signal connection closed at TCP layer

Figure 5-8 shows a trace of timestamp difference accepted.

**Figure 5-8 Trace of Timestamp Difference Accepted**

| No. | Time     | Source     | Destination | Protocol | Info                                                       |
|-----|----------|------------|-------------|----------|------------------------------------------------------------|
| 27  | 0.950015 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 > 64148 [ACK] Seq=3824906691 Ack=3324124303 wln=32768 |
| 28  | 0.950478 | ff.ff.fd   | ff.ff.fd    | SW_ILS   | ELP                                                        |
| 29  | 0.950153 | 10.10.11.2 | 10.10.10.2  | TCP      | 3225 > 64148 [ACK] Seq=3824906691 Ack=3324124471 wln=32768 |
| 30  | 0.950177 | ff.ff.fd   | ff.ff.fd    | FC       | Link CTL, ACK                                              |
| 31  | 0.950014 | 10.10.10.2 | 10.10.11.2  | TCP      | 64148 > 3225 [ACK] Seq=3324124471 Ack=3824906755 wln=32768 |
| 32  | 0.950990 | ff.ff.fd   | ff.ff.fd    | SW_ILS   | SW_ACC (ELP)                                               |
| 33  | 0.950014 | 10.10.10.2 | 10.10.11.2  | TCP      | 64148 > 3225 [ACK] Seq=3324124471 Ack=3824906823 wln=32768 |

|                                                       |  |
|-------------------------------------------------------|--|
| FCIP (SOFF/BOFN)                                      |  |
| Protocol: 1                                           |  |
| Version: 1                                            |  |
| Protocol (L's Complement): 254                        |  |
| Version (L's Complement): 254                         |  |
| FCIP Encapsulation word: 0x0101fefe                   |  |
| . = changed flag: false                               |  |
| 0 = special frame flag: false                         |  |
| PFlags (L's Complement): 0xff                         |  |
| 0000 00.. = Flags: 0x00                               |  |
| ....00 0010 1010 = Frame Length (L's Complement): 42  |  |
| 1111 11.. = Flags (L's Complement): 0x3f              |  |
| ....11 1101 0101 = Frame Length (L's Complement): 981 |  |
| Time (secs): 3.04755486                               |  |
| Time (fraction): 3654381726                           |  |
| CRC: 0x00000000                                       |  |
| SOFF: SOFF (0x28)                                     |  |
| SOFF (L's Complement): 0xd7                           |  |
| BOFN: BOFN (0x41)                                     |  |
| BOFN (L's Complement): 0xb6                           |  |
| More channel                                          |  |
| R_CTL: 0x02                                           |  |

Timestamp used to check acceptable time difference

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

## FCIP Special Frame Tunnel Creation and Monitoring

Previous FCIP tunnel configuration must be completed before adding FCIP Special Frame configuration. This section describes how to correctly configure and show an FCIP tunnel with a Special Frame.

```
MDS2# show wwn switch
```

```
Switch WWN is 20:00:00:0b:5f:d5:9f:c0
```

(You'll need the WWN of each MDS 9000 switch end point.)

```
MDS1(config)# int fcip 21
```

```
MDS1(config-if)# special-frame peer-wwn 20:00:00:0b:5f:d5:9f:c0
```

(This enables the Special Frame that is used in the creation of the FCIP tunnel.)

```
MDS1# show wwn switch
```

```
Switch WWN is 20:00:00:05:30:00:59:de
```

```
MDS2(config)# int fcip 21
```

```
MDS2(config-if)# special-frame peer-wwn 20:00:00:05:30:00:59:de
```

```
module-2#
```

```
Jan 14 15:25:38 port1: 857314:FCIP21: SUP-> Set Port mode 1
```

```
Jan 14 15:25:38 port1: 857315:FCIP21: SUP-> Port VSAN (1) already set to same value
```

```
Jan 14 15:25:38 port1: 857316:FCIP21: SUP-> Trunk mode (1) already set to same value
```

```
Jan 14 15:25:38 port1: 857317:FCIP21: SUP-> Enable tunnel ADMIN UP
```

```
Jan 14 15:25:38 port1: 857318:FCIP21: Try to Bring UP the Tunnel
```

```
Jan 14 15:25:38 port1: 857319:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
```

```
Jan 14 15:25:38 port1: 857320:FCIP: Create a new listener object for 10.10.11.2:3225
```

```
Jan 14 15:25:38 port1: 857321:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
```

```
Jan 14 15:25:38 port1: 857322:FCIP21: Create a DE 0xd802cd00 for this tunnel
```

```
Jan 14 15:25:38 port1: 857323:FCIP21: Bind the DE 0xd802cd00 [1] to tunnel LEP 0x80111570
```

```
Jan 14 15:25:38 port1: 857324:FCIP21: Start the active connection [1] to 10.10.10.2:3225
```

```
Jan 14 15:25:38 port1: 857325:FCIP21: Create a DE 0xd802db40 for this tunnel
```

```
Jan 14 15:25:38 port1: 857326:FCIP21: Bind the DE 0xd802db40 [2] to tunnel LEP 0x80111570
```

```
Jan 14 15:25:38 port1: 857327:FCIP21: Start the active connection [2] to 10.10.10.2:3225
```

```
Jan 14 15:25:38 port1: 857328:FCIP21: Active Connect creation SUCCEEDED [1]
```

```
Jan 14 15:25:38 port1: 857329:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
```

```
Jan 14 15:25:38 port1: 857330:FCIP21: Setup for Special Frame handling: I'm Originator
```

(This begins the Special Frame setup of the Originator.)

```
Jan 14 15:25:38 port1: 857331:FCIP21: Send the SF as Originator & wait for response
```

(The Special Frame is sent.)

```
Jan 14 15:25:38 port1: 857332:FCIP21: Setup timer to wait for SF
```

```
Jan 14 15:25:38 port1: 857333:FCIP21: Active Connect creation SUCCEEDED [2]
```

(The Special Frame is correctly configured with the WWN of the remote MDS 9000 switch.)

```
Jan 14 15:25:38 port1: 857334:FCIP21: Bind DE 2 to TCP-hdl 0xd8072000
```

```
Jan 14 15:25:38 port1: 857335:FCIP21: Setup for Special Frame handling: I'm Originator
```

```
Jan 14 15:25:38 port1: 857336:FCIP21: Send the SF as Originator & wait for response
```

```
Jan 14 15:25:38 port1: 857337:FCIP21: Setup timer to wait for SF
```

```
Jan 14 15:25:38 port1: 857338:FCIP21: processing SF frame, I'm Originator
```

```
Jan 14 15:25:38 port1: 857339:FCIP21: Bind DE 1 to eport 0x80110550
```

```
Jan 14 15:25:38 port1: 857340:FCIP21: bind de 1 in eport 0x80110550, hash = 1 num-conn: 2
```

```
Jan 14 15:25:38 port1: 857341:FCIP21: processing SF frame, I'm Originator
```

```
Jan 14 15:25:38 port1: 857342:FCIP21: Bind DE 2 to eport 0x80110550
```

```
Jan 14 15:25:38 port1: 857343:FCIP21: bind de 2 in eport 0x80110550, hash = 2 num-conn: 2
```

```
Jan 14 15:25:38 port1: 857344:FCIP21: Send LINK UP to SUP
```

```
Jan 14 15:25:39 port1: 857345:FCIP21: SUP-> Set trunk mode: 2
```

```
Jan 14 15:25:39 port1: 857346:FCIP21: Change the operational mode to TRUNK
```

```
Jan 14 15:25:39 port1: 857347:FCIP21: *** Received non-eisl frame in TE mode 64 64
```

```
MDS2# show int fcip 21
```

```
fcip21 is trunking
```



*Send comments to mdsfeedback-doc@cisco.com.*

```
Hardware is GigabitEthernet
Port WWN is 20:42:00:0b:5f:d5:9f:c0
Peer port WWN is 20:42:00:05:30:00:59:de
Admin port mode is auto, trunk mode is on
Port mode is TE
vsan is 1
Trunk vsans (allowed active) (1-2)
Trunk vsans (operational) (1-2)
Trunk vsans (up) (1-2)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
Using Profile id 21 (interface GigabitEthernet2/1)
Peer Information
 Peer Internet address is 10.10.10.2 and port is 3225
 Special Frame is enabled
```

(The Special Frame is enabled. It is used for security to verify that the tunnel remote end point is the correct pwwn of the switch.)

```
Peer switch WWN is 20:00:00:05:30:00:59:de
```

(This is the peer WWN of the remote switch. The pWWN of the switch can be found using the **show wwn switch** command.)

```
Maximum number of TCP connections is 2
Time Stamp is enabled, acceptable time difference 3000 ms
B-port mode disabled
TCP Connection Information
 2 Active TCP connections
 Control connection: Local 10.10.11.2:64792, Remote 10.10.10.2:3225
 Data connection: Local 10.10.11.2:64794, Remote 10.10.10.2:3225
 372 Attempts for active connections, 345 close of connections
TCP Parameters
 Path MTU 1500 bytes
 Current retransmission timeout is 300 ms
 Round trip time: Smoothed 10 ms, Variance: 5
 Advertized window: Current: 64 KB, Maximum: 64 KB, Scale: 1
 Peer receive window: Current: 64 KB, Maximum: 64 KB, Scale: 1
 Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
```

*Send comments to mdsfeedback-doc@cisco.com.*

Figure 5-9 shows a trace of an FCIP tunnel with a Special Frame.

**Figure 5-9 Trace of FCIP Tunnel with a Special Frame**

| No. | Time     | Source     | Destination | Protocol | Info                                                       |
|-----|----------|------------|-------------|----------|------------------------------------------------------------|
| 10  | 2.964705 | 10.10.11.2 | 10.10.10.2  | TCP      | 64790 > 3225 [ACK] Seq=2988533241 Ack=3249656006 win=32768 |
| 11  | 2.964778 | 10.10.11.2 | 10.10.10.2  | FCIP     | Special Frame                                              |
| 12  | 2.964791 | 10.10.11.2 | 10.10.10.2  | TCP      | 64788 > 3225 [ACK] Seq=2988533241 Ack=3249656006 win=32768 |
| 13  | 2.964820 | 10.10.10.2 | 10.10.11.2  | FCIP     | Special Frame                                              |
| 14  | 2.964824 | 10.10.10.2 | 10.10.11.2  | TCP      | 3225 > 64790 [ACK] Seq=3230217584 Ack=2937578959 win=32768 |
| 15  | 2.964837 | 10.10.10.2 | 10.10.11.2  | TCP      | 3225 > 64790 [ACK] Seq=3230217584 Ack=2937578959 win=32768 |
| 16  | 2.964850 | 10.10.10.2 | 10.10.11.2  | FCIP     | Special Frame                                              |
| 17  | 2.964867 | 10.10.10.2 | 10.10.11.2  | TCP      | 3225 > 64788 [ACK] Seq=3249656006 Ack=2968533241 win=32768 |
| 18  | 2.964881 | 10.10.10.2 | 10.10.11.2  | TCP      | 3225 > 64788 [ACK] Seq=3249656006 Ack=2968533317 win=32768 |

```

Transmission Control Protocol, Src Port: 64790 (64790), Dst Port: 3225 (3225), Seq: 2937578959, Ack: 3230217584, Len: 76
FCIP
 Protocol: 3 Protocol and Version always one
 Version: 1
 Protocol (1's Complement): 254
 Version (1's Complement): 254 One complement of above 1
 FCIP Encapsulation word: 0x0101fafe Previous four bytes repeated
 - = Changed flag: False
 1 = Special Frame Flag: True Special Frame bit enabled
 rFlags (1's complement): 0xfe1's Complement of Special Frame Flag: True
 0000 00.. = Flags: 0x00
 00 0001 0011 = Frame Length (in words): 10
 1111 13.. = Flags (1's Complement): 0x2f
 11 1110 1100 = Frame Length (1's Complement): 1004
 Time (secs): 1042558283
 Time (fraction): 1323647828
 CRC: 0x00800960
 Source Fabric WWN: 20:08:00:0b:5f:d5:9f:c0 (00:0b:5f)
 FC/FCIP Entity ID: 0000000000000015 "fcip profile 21" used on MDS configuration. Hex 15 = Dec 21
 Connection Name: 0000000000000000
 Connection Usage Flags: 0x00
 Connection Usage Code: 0x0080
 Destination Fabric WWN: 00:05:10:08:19:da:00:00 wwn of remote MDS switch
 K_A_Tov: 0

```

## Special Frame Misconfiguration Examples

The following example shows an incorrect peer WWN when using Special Frame.

```

module-2# Jan 14 15:14:30 port1: 855278:FCIP21: SUP-> Set Port mode 1
Jan 14 15:14:30 port1: 855279:FCIP21: SUP-> Port VSAN (1) already set to same value
Jan 14 15:14:30 port1: 855280:FCIP21: SUP-> Trunk mode (1) already set to same
Jan 14 15:14:30 port1: 855281:FCIP21: SUP-> Enable tunnel ADMIN UP
Jan 14 15:14:30 port1: 855282:FCIP21: Try to Bring UP the Tunnel
Jan 14 15:14:30 port1: 855283:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Jan 14 15:14:30 port1: 855284:FCIP: Create a new listener object for 10.10.11.2:3225
Jan 14 15:14:30 port1: 855285:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Jan 14 15:14:30 port1: 855286:FCIP21: Create a DE 0xd802d240 for this tunnel
Jan 14 15:14:30 port1: 855287:FCIP21: Bind the DE 0xd802d240 [1] to tunnel LEP 0x80111570
Jan 14 15:14:30 port1: 855288:FCIP21: Start the active connection [1] to 10.10.10.2:3225
Jan 14 15:14:30 port1: 855289:FCIP21: Create a DE 0xd802d200 for this tunnel
Jan 14 15:14:30 port1: 855290:FCIP21: Bind the DE 0xd802d200 [2] to tunnel LEP 0x80111570
Jan 14 15:14:30 port1: 855291:FCIP21: Start the active connection [2] to 10.10.10.2:3225
Jan 14 15:14:30 port1: 855292:FCIP21: Active Connect creation SUCCEEDED [1]
Jan 14 15:14:30 port1: 855293:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
Jan 14 15:14:30 port1: 855294:FCIP21: Setup for Special Frame handling: I'm Originator
Jan 14 15:14:30 port1: 855295:FCIP21: Send the SF as Originator & wait for response
Jan 14 15:14:30 port1: 855296:FCIP21: Setup timer to wait for SF
Jan 14 15:14:30 port1: 855297:FCIP21: Active Connect creation SUCCEEDED [2]
Jan 14 15:14:30 port1: 855298:FCIP21: Bind DE 2 to TCP-hdl 0xd8072000
Jan 14 15:14:30 port1: 855299:FCIP21: Setup for Special Frame handling: I'm Originator
Jan 14 15:14:30 port1: 855300:FCIP21: Send the SF as Originator & wait for response
Jan 14 15:14:30 port1: 855301:FCIP21: Setup timer to wait for SF
Jan 14 15:14:30 port1: 855302:FCIP21: TCP Received a close connection [1] reason 1
Jan 14 15:14:30 port1: 855303:FCIP21: Delete the DE [1]0xd802d240

```

*Send comments to mdsfeedback-doc@cisco.com.*

```

Jan 14 15:14:30 port1: 855304:FCIP21: DE [-670903744] 0x00000001 terminate tcp connection
0xd8072c00
Jan 14 15:14:30 port1: 855305:FCIP21: Delete the DE object [1] 0xd802d240
Jan 14 15:14:30 port1: 855306:FCIP21: lep not bound, close only de [1]
Jan 14 15:14:30 port1: 855307:FCIP21: TCP Received a close connection [2] reason 1
Jan 14 15:14:30 port1: 855308:FCIP21: Delete the DE [2]0xd802d200
Jan 14 15:14:30 port1: 855309:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:30 port1: 855310:FCIP21: Start the bringup tunnel timer, timeout: 38740
Jan 14 15:14:30 port1: 855311:FCIP21: DE [-670903808] 0x00000002 terminate tcp connection
0xd8072000
Jan 14 15:14:30 port1: 855312:FCIP21: Delete the DE object [2] 0xd802d200
Jan 14 15:14:30 port1: 855313:FCIP21: lep not bound, close only de [2]
Jan 14 15:14:31 port1: 855314:FCIP21: Received new TCP connection from peer:
10.10.10.2:64050
Jan 14 15:14:31 port1: 855315:FCIP21: Create a DE 0xd802d080 for this tunnel
Jan 14 15:14:31 port1: 855316:FCIP21: Bind the DE 0xd802d080 [1] to tunnel LEP 0x80111570
Jan 14 15:14:31 port1: 855317:FCIP21: Bind DE 1 to TCP-hdl 0xd8072000
Jan 14 15:14:31 port1: 855318:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:31 port1: 855319:FCIP21: Setup timer to wait for SF
Jan 14 15:14:31 port1: 855320:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:31 port1: 855321:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:31 port1: 855322:FCIP21: Delete the DE [1]0xd802d080
Jan 14 15:14:31 port1: 855323:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:31 port1: 855324:FCIP21: DE [-670904192] 0x00000001 terminate tcp connection
0xd8072000
Jan 14 15:14:31 port1: 855325:FCIP21: Delete the DE object [1] 0xd802d080
Jan 14 15:14:31 port1: 855326:FCIP21: Received new TCP connection from peer:
10.10.10.2:64048
Jan 14 15:14:31 port1: 855327:FCIP21: Create a DE 0xd802d200 for this tunnel
Jan 14 15:14:31 port1: 855328:FCIP21: Bind the DE 0xd802d200 [1] to tunnel LEP 0x80111570
Jan 14 15:14:31 port1: 855329:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
Jan 14 15:14:31 port1: 855330:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:31 port1: 855331:FCIP21: Setup timer to wait for SF
Jan 14 15:14:31 port1: 855332:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:31 port1: 855333:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:31 port1: 855334:FCIP21: Delete the DE [1]0xd802d200
Jan 14 15:14:31 port1: 855335:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:31 port1: 855336:FCIP21: DE [-670903808] 0x00000001 terminate tcp connection
0xd8072c00
Jan 14 15:14:31 port1: 855337:FCIP21: Delete the DE object [1] 0xd802d200
Jan 14 15:14:37 port1: 855338:FCIP21: Received new TCP connection from peer:
10.10.10.2:64046
Jan 14 15:14:37 port1: 855339:FCIP21: Create a DE 0xd802d5c0 for this tunnel
Jan 14 15:14:37 port1: 855340:FCIP21: Bind the DE 0xd802d5c0 [1] to tunnel LEP 0x80111570
Jan 14 15:14:37 port1: 855341:FCIP21: Bind DE 1 to TCP-hdl 0xd8071000
Jan 14 15:14:37 port1: 855342:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:37 port1: 855343:FCIP21: Setup timer to wait for SF
Jan 14 15:14:37 port1: 855344:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:37 port1: 855345:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:37 port1: 855346:FCIP21: Delete the DE [1]0xd802d5c0
Jan 14 15:14:37 port1: 855347:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:37 port1: 855348:FCIP21: DE [-670902848] 0x00000001 terminate tcp connection
0xd8071000
Jan 14 15:14:37 port1: 855349:FCIP21: Delete the DE object [1] 0xd802d5c0
Jan 14 15:14:37 port1: 855350:FCIP21: Received new TCP connection from peer:
10.10.10.2:64044
Jan 14 15:14:37 port1: 855351:FCIP21: Create a DE 0xd802cac0 for this tunnel
Jan 14 15:14:37 port1: 855352:FCIP21: Bind the DE 0xd802cac0 [1] to tunnel LEP 0x80111570
Jan 14 15:14:37 port1: 855353:FCIP21: Bind DE 1 to TCP-hdl 0xd8071400
Jan 14 15:14:37 port1: 855354:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:37 port1: 855355:FCIP21: Setup timer to wait for SF

```

*Send comments to mdsfeedback-doc@cisco.com.*

```

Jan 14 15:14:37 port1: 855356:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:37 port1: 855357:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:37 port1: 855358:FCIP21: Delete the DE [1]0xd802cac0
Jan 14 15:14:37 port1: 855359:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:37 port1: 855360:FCIP21: DE [-670905664] 0x00000001 terminate tcp connection
0xd8071400
Jan 14 15:14:37 port1: 855361:FCIP21: Delete the DE object [1] 0xd802cac0

```

Figure 5-10 shows a trace of an incorrect remote switch WWN using a Special Frame

**Figure 5-10 Trace of Incorrect Remote Switch WWN Using a Special Frame**

| No. | Time     | Source     | Destination | Protocol | Info                                                       |
|-----|----------|------------|-------------|----------|------------------------------------------------------------|
| 10  | 2.964755 | 10.10.10.2 | 10.10.10.2  | FCIP     | Special Frame                                              |
| 11  | 2.964779 | 10.10.10.2 | 10.10.10.2  | TCP      | 64788 > 3225 [ACK] Seq=2968533241 Ack=3249656006 win=32768 |
| 12  | 2.964791 | 10.10.10.2 | 10.10.10.2  | FCIP     | Special Frame                                              |
| 13  | 2.964810 | 10.10.10.2 | 10.10.10.2  | TCP      | 3225 > 64790 [ACK] Seq=3230217584 Ack=2937578959 win=32768 |
| 14  | 2.964824 | 10.10.10.2 | 10.10.10.2  | TCP      | 3225 > 64790 [ACK] Seq=3230217584 Ack=2937578959 win=32768 |
| 15  | 2.964837 | 10.10.10.2 | 10.10.10.2  | FCIP     | Special Frame                                              |
| 16  | 2.964850 | 10.10.10.2 | 10.10.10.2  | TCP      | 3225 > 64788 [ACK] Seq=3249656006 Ack=2968533241 win=32768 |
| 17  | 2.964867 | 10.10.10.2 | 10.10.10.2  | TCP      | 3225 > 64788 [ACK] Seq=3249656006 Ack=2968533241 win=32768 |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| Internet Protocol, Src Addr: 10.10.10.2 (10.10.10.2), Dst Addr: 10.10.10.2 (10.10.10.2)<br>Transmission Control Protocol, Src Port: 64790 (64790), Dst Port: 3225 (3225), Seq: 2937578959, Ack: 3230217584, Len: 76<br>FCIP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |  |
| Protocol: 3 Protocol and Version always use<br>Version: 1<br>Protocol (1's complement): 254 One complement of above 1<br>Version (1's complement): 254<br>FCIP encapsulation word: 0x0101fa Previous four bytes repeated<br>- = changed flag: false<br>1 = Special Frame Flag: True Special Frame bit enabled<br>Rflags (1's complement): 0xfef1's Complement of Special Frame Flag: True<br>0000 00.. = Flags: 0x00<br>.... 0001 0011 = Frame Length (in words): 19<br>1111 11.. = Flags (1's complement): 0x3f<br>.... 1111 1100 = Frame Length (1's complement): 1004<br>time (secs): 1042558283<br>time (fraction): 1323647828<br>CRC: 0x00000000<br>Source Fabric WWN: 20:08:00:0b:5f:d5:9f:c0 (00:0b:5f)<br>FC/FCIP Entity ID: 0000000000000015 "tcp profile 21" used on MDS configuration. Hex 15 = Dec 21<br>Connection Name: 00000000EAD8EEF<br>Connection Usage Flags: 0x00<br>Connection Usage Code: 0x0000<br>Destination Fabric WWN: 00:15:10:08:59:de:00:00 wwn of remote MDS switch<br>K_A_Tov: 0 |  |

*Send comments to mdsfeedback-doc@cisco.com.*

## Troubleshooting iSCSI Issues

There are several types of issues you can experience with iSCSI, including the following:


- Troubleshooting iSCSI Authentication, page 5-51
- Configuring Authentication, page 5-52
- Troubleshooting Username/Password Configuration, page 5-53
- Troubleshooting Radius Configuration, page 5-53
- Troubleshooting Radius Routing Configuration, page 5-56
- Troubleshooting Dynamic iSCSI Configuration, page 5-56

## Troubleshooting iSCSI Authentication

iSCSI user login authentication is required with the Cisco MDS 9000 Family switch. There are two ways of getting iSCSI users authenticated: either locally configured in the switch's configuration file, or using the Radius server database.

Figure 5-11 shows a successful iSCSI login for the Windows 2000 driver.

**Figure 5-11 Successful iSCSI Login Status Window**




On Solaris systems, a successful login is found in the /var/adm/messages directory, and should look similar to the following example:

```
Mar 14 12:53:23 ca-sun1 iscsid[12745]: [ID 702911 daemon.notice] discovery process for
172.22.91.223 finished, exiting
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 448557 daemon.notice] logged into
DiscoveryAddress 172.22.91.223:3260 isid 023d0040
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 702911 daemon.notice] iSCSI target 2 =
iqn.com.domainname.vrrp-11.gw.21000020375aff77 at0
Mar 14 12:58:45 ca-sun1 iscsid[12809]: [ID 529321 daemon.notice] logged into target
iqn.com.domainname.vrrp-11.gw.21000020375aff77 7
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 702911 daemon.notice] iSCSI target 3 =
iqn.com.domainname.vrrp-11.gw.21000020374baf02 at0
Mar 14 12:58:45 ca-sun1 iscsid[12810]: [ID 529321 daemon.notice] logged into target
iqn.com.domainname.vrrp-11.gw.21000020374baf02 7
```

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

Figure 5-12 shows a failed iSCSI login for the Windows 2000 driver.

**Figure 5-12 Failed iSCSI Login Status Window**



On Solaris systems, a failed login is found in the `/var/adm/messages` directory and should look similar to the following example.

```
Mar 14 11:44:42 ca-sun1 iscsid[12561]: [ID 702911 daemon.notice] login rejected: initiator
error (01)
Mar 14 11:44:42 ca-sun1 iscsid[12561]: [ID 702911 daemon.error] Hard discovery login
failure to 172.22.91.223:3260 - exiting
Mar 14 11:44:42 ca-sun1 iscsid[12561]: [ID 702911 daemon.notice] discovery process for
172.22.91.223 finished, exiting
```

## Configuring Authentication

Whenever you experience a login failure, use the **show authentication** command to see if the iSCSI authentication is correctly defined. A sample of local authentication should look like this:

[illegible]

If iSCSI is configured for radius authentication, it should look like this:

[illegible]

```
telnet/ssh: not enabled
iscsi: enabled
authentication method: local
console: enabled
telnet/ssh: enabled
iscsi: enabled
switch#
```

The client side username and password should be check against either the switch's local configuration file or the Radius user database.

```
switch# sh user-account iscsi
username:iscsi
secret:1234567812345678

username:iscsiuser
secret:1234567812345678
```

If authentication is against the Radius server, ping the Radius server to and from the switch to make sure it can be reached over IP. Execute the **show radius-server** command to make sure radius key and port for authentication and accounting match exactly with is configured on Radius server.

```
switch# show radius-server
retransmission count:3
timeout value:5

following RADIUS servers are configured:
 171.71.49.197:
 available for authentication on port:1812
 available for accounting on port:1813
 RADIUS shared secret:radius
```

Adjust the radius timeout and retransmission accordingly, as they have default value of 1 sec and 1 time. Figure 5-13 shows a Windows-based Radius server configuration.

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

**Figure 5-13 Windows-Based Raduis Server Configuration Dialog**



If the items shown above match, verify that the client username and password match those in the Radius database.

The following example shows the results of the **debug security radius** command, if the iSCSI client logs in successfully.

```
switch#
switch# Mar 4 23:16:20 securityd: received CHAP authentication request for user002
Mar 4 23:16:20 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar 4 23:16:20 securityd: reading RADIUS configuration
Mar 4 23:16:20 securityd: opening radius configuration for group:default
Mar 4 23:16:20 securityd: opened the configuration successfully
Mar 4 23:16:20 securityd: GET request for RADIUS global config
Mar 4 23:16:20 securityd: got back the return value of global radius configuration
operation:success
Mar 4 23:16:20 securityd: closing RADIUS pss configuration
Mar 4 23:16:20 securityd: opening radius configuration for group:default
Mar 4 23:16:20 securityd: opened the configuration successfully
Mar 4 23:16:20 securityd: GETNEXT request for radius index:0 addr:
Mar 4 23:16:20 securityd: got some reply from 171.71.49.197
Mar 4 23:16:20 securityd: verified the response from:171.71.49.197
Mar 4 23:16:20 securityd: RADIUS server sent accept for authentication request for
user002
Mar 4 23:16:25 securityd: received CHAP authentication request for user002
Mar 4 23:16:25 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar 4 23:16:25 securityd: reading RADIUS configuration
Mar 4 23:16:25 securityd: opening radius configuration for group:default
Mar 4 23:16:25 securityd: opened the configuration successfully
Mar 4 23:16:25 securityd: GET request for RADIUS global config
```



*Send comments to mdsfeedback-doc@cisco.com.*

```
Mar 4 23:16:25 securityd: got back the return value of global radius configuration
operation:success
Mar 4 23:16:25 securityd: closing RADIUS pss configuration
Mar 4 23:16:25 securityd: opening radius configuration for group:default
Mar 4 23:16:25 securityd: opened the configuration successfully
Mar 4 23:16:25 securityd: GETNEXT request for radius index:0 addr:
Mar 4 23:16:25 securityd: got some reply from 171.71.49.197
Mar 4 23:16:25 securityd: verified the response from:171.71.49.197
Mar 4 23:16:25 securityd: RADIUS server sent accept for authentication request for
user002
Mar 4 23:16:25 securityd: got some reply from 171.71.49.197
Mar 4 23:16:25 securityd: verified the response from:171.71.49.197
Mar 4 23:16:25 securityd: RADIUS server sent accept for authentication request for
user002
```

The example above shows that the iSCSI client has been authenticated 3 times, first for the switch login, and the second and third times for the SCSI drive login. The switch sends Radius attributes 1, 3, 4, 5, 6, 60 and 61 to the Radius server. The Radius server only needs to respond with **request accept** or **request reject**.

The following example shows a radius authentication.

```
639 2003y3m14d 15h12m48s -----
640 2003y3m14d 15h12m48s Message Type=Access_Request
641 2003y3m14d 15h12m48s ID=243, Length=90
642 2003y3m14d 15h12m48s User name=user002
643 2003y3m14d 15h12m48s NAS IP address=2887147911
644 2003y3m14d 15h12m48s CHAP password=%j+<.Wøøë-K-ëÛ<]
645 2003y3m14d 15h12m48s CHAP challenge=n8NŸgø$"__ô4}ôx
646 2003y3m14d 15h12m48s NAS port=1426
647 2003y3m14d 15h12m48s NAS port type=5
648 2003y3m14d 15h12m48s Service type=8
649 2003y3m14d 15h12m48s User (user002) authenticate OK.
650 2003y3m14d 15h12m54s -----
651 2003y3m14d 15h12m54s Message Type=Access_Request
652 2003y3m14d 15h12m54s ID=60, Length=90
653 2003y3m14d 15h12m54s User name=user002
654 2003y3m14d 15h12m54s NAS IP address=2887147911
655 2003y3m14d 15h12m54s CHAP password=_çĒò_à!_AëC0__`ö
656 2003y3m14d 15h12m54s CHAP challenge=_/ô½Ÿ×!âëË 4_`ZH
657 2003y3m14d 15h12m54s NAS port=1426
658 2003y3m14d 15h12m54s NAS port type=5
659 2003y3m14d 15h12m54s Service type=8
660 2003y3m14d 15h12m54s User (user002) authenticate OK.
661 2003y3m14d 15h12m54s -----
662 2003y3m14d 15h12m54s Message Type=Access_Request
663 2003y3m14d 15h12m54s ID=179, Length=90
664 2003y3m14d 15h12m54s User name=user002
665 2003y3m14d 15h12m54s NAS IP address=2887147911
666 2003y3m14d 15h12m54s CHAP password=--5Âùrfâxh
667 2003y3m14d 15h12m54s CHAP challenge=#ùĒŸü{"__"´_Ux
668 2003y3m14d 15h12m54s NAS port=1426
669 2003y3m14d 15h12m54s NAS port type=5
670 2003y3m14d 15h12m54s Service type=8
671 2003y3m14d 15h12m54s User (user002) authenticate OK.
```

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

## Troubleshooting Radius Routing Configuration

The switch sends the Radius authentication request from the mgmt0 interface, so the correct route to the Radius server must be defined. If no correct route is defined, the switch may send the Radius request from Gigabit Ethernet port. In that case, the Radius server returns the accept to the Gigabit Ethernet port and the switch does not get the response. The following example shows the output from the **debug security radius** command.

```
switch# Mar 5 00:51:13 securityd: received CHAP authentication request for user002
Mar 5 00:51:13 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar 5 00:51:13 securityd: reading RADIUS configuration
Mar 5 00:51:13 securityd: opening radius configuration for group:default
Mar 5 00:51:13 securityd: opened the configuration successfully
Mar 5 00:51:13 securityd: GET request for RADIUS global config
Mar 5 00:51:13 securityd: got back the return value of global radius configuration
operation:success
Mar 5 00:51:13 securityd: closing RADIUS pss configuration
Mar 5 00:51:13 securityd: opening radius configuration for group:default
Mar 5 00:51:13 securityd: opened the configuration successfully
Mar 5 00:51:13 securityd: GETNEXT request for radius index:0 addr:
Mar 5 00:51:18 securityd: sending data to 171.71.49.197
Mar 5 00:51:18 securityd: waiting for response from 171.71.49.197
Mar 5 00:51:23 securityd: sending data to 171.71.49.197
Mar 5 00:51:23 securityd: waiting for response from 171.71.49.197
Mar 5 00:51:28 securityd: sending data to 171.71.49.197
Mar 5 00:51:28 securityd: waiting for response from 171.71.49.197
Mar 5 00:51:33 securityd: trying out next server
Mar 5 00:51:33 securityd: no response from RADIUS server for authentication user002
Mar 5 00:51:33 securityd: doing local chap authentication for user002
Mar 5 00:51:33 securityd: local chap authentication result for user002:user not present
```

## Troubleshooting Dynamic iSCSI Configuration

A physical Fibre Channel target (target pWWN) presented as an iSCSI target, makes the physical targets accessible to iSCSI hosts. The IPS module presents physical Fibre Channel targets as iSCSI targets to iSCSI hosts in one of two ways: Dynamic Mapping or Static Mapping.

By default, the IPS module does not automatically import Fibre Channel targets. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have the configured name. Targets that are not mapped will be advertised with the name created by the conventions explained in this section.

### Checking the Configuration

Verify the configuration of the Gigabit Ethernet Interface by performing the following steps.

- Ensure that you are configuring the proper slot or port.
- Ensure that the Gigabit Ethernet interfaces are not shut down. Each Gigabit Ethernet interface is “partnered” with a virtual iSCSI interface. In order for iSCSI to operate on a particular Gigabit Ethernet, the virtual iSCSI interface for that port must be in a “no shutdown” state:

```
interface Gigabit Ethernet 3/1
no shutdown
.
```

*Send comments to mdsfeedback-doc@cisco.com.*

```
.
.
interface iscsi 3/1
no shutdown
```

- Verify that the IP parameters are correct.
- Verify authentication on gigabit Ethernet interface (None or Chap) matches the authentication configured on the iSCSI initiator. Note that configuring authentication at the interface level overrides the Global Authentication setting.
- Verify gigabit Ethernet switchport parameters are correct (MTU, mode, etc.).

## Performing Basic Dynamic iSCSI Troubleshooting

Keep the following in mind when performing basic dynamic iSCSI troubleshooting:

- **iscsi import target fc** must be enabled in order to allow SCSI targets to be discovered by the logged-in iSCSI initiators.
- Dynamic iSCSI configuration places all iSCSI initiators logging into the MDS9000 into VSAN 1 by default.
- Any zoning in effect on the default VSAN (VSAN1) will also be applied to iSCSI-connected devices.

## Useful `show` Commands for Debugging Dynamic iSCSI Configuration

The output from the following commands reflects correctly established iSCSI sessions. Execute the same commands on your switch and compare with the output below to help identify possible issues:

**show iscsi session detail**

**show iscsi remote-node initiator**

**show iscsi stats**

**show iscsi stats detail**

**show iscsi local-node**

**show fcns data vsan 1**

**show flogi database vsan 1**

### show iscsi session detail

```
switch#show iscsi session detail
Initiator ign.1987-05.com.cisco.02.F984BCA7E08C307E2D87A099B2D452F3.FULLMOON (FULLMOON)
Session #1 (index 2)
 Target ign.com.domainname.IPS-TEST.02-07.gw.202300a0b80b14da
 VSAN 1, ISID 000000000000, TSID 134, Status active, no reservation
 Type Normal, ExpCmdSN 44, MaxCmdSN 53, Barrier 0
 MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
 DataSeqInOrder No, InitialR2T Yes, ImmediateData No
 Registered LUN 0, Mapped LUN 0
 Stats:
 PDU: Command: 42, Response: 36
```

*Send comments to mdsfeedback-doc@cisco.com.*

```

Bytes: TX: 4960, RX: 0
Number of connection: 1
Connection #1
 Local IP address: 0xa021ec8, Peer IP address: 0xa021eca
 CID 0, State: LOGGED_IN
 StatSN 43, ExpStatSN 0
 MaxRecvDSLength 524288, our_MaxRecvDSLength 1024
 CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
 AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
 Version Min: 0, Max: 0
 FC target: Up, Reorder PDU: No, Marker send: No (int 0)
 Received MaxRecvDSLen key: Yes

```

### show iscsi remote-node initiator

```

switch# sh iscsi remote-node initiator
iSCSI Node name is ign.1987-05.com.cisco.02.F984BCA7E08C307E2D87A099B2D452F3.FULLMOON
 iSCSI alias name: FULLMOON
 Node WWN is 20:0c:00:0b:be:77:72:42 (dynamic)
 Member of vsans: 1
 Number of Virtual n_ports: 1
 Virtual Port WWN is 20:0d:00:0b:be:77:72:42 (dynamic)
 Interface iSCSI 2/7, Portal group tag: 0x86
 VSAN ID 1, FCID 0x750105

```

### show iscsi local-node

```

switch# sh iscsi local-node
target: ign.com.domainname.IPS-TEST.02-07.gw.202300a0b80b14da
 Port WWN 20:23:00:a0:b8:0b:14:da , VSAN 1
 Auto-created node

```

### show fcns data vsan 1

```

switch# sh fcns data vsan 1

VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x750000 N 20:23:00:a0:b8:0b:14:da (SymBios) scsi-fcp:target
0x750102 N 10:00:00:00:c9:30:ba:06 (Emulex) scsi-fcp:init
0x750105 N 20:0d:00:0b:be:77:72:42 scsi-fcp:init isc..w
0x750201 N 50:08:05:f3:00:04:96:71 scsi-fcp
0x750301 N 50:08:05:f3:00:04:96:79 scsi-fcp
0x750400 N 20:00:00:02:3d:07:05:c0 (NuSpeed) scsi-fcp:init

```

### show flogi databse vsan 1

```

switch# show flogi database vsan 1

INTERFACE VSAN FCID PORT NAME NODE NAME

fc1/1 1 0x750400 20:00:00:02:3d:07:05:c0 10:00:00:02:3d:07:05:c0
fc1/6 1 0x750000 20:23:00:a0:b8:0b:14:da 20:22:00:a0:b8:0b:14:d9
fc1/8 1 0x750102 10:00:00:00:c9:30:ba:06 20:00:00:00:c9:30:ba:06
fc1/9 1 0x750201 50:08:05:f3:00:04:96:71 50:08:05:f3:00:04:96:70
fc1/10 1 0x750301 50:08:05:f3:00:04:96:79 50:08:05:f3:00:04:96:70
iscsi2/7 1 0x750105 20:0d:00:0b:be:77:72:42 20:0c:00:0b:be:77:72:42

```

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

## Virtual Target Access Control

When creating a virtual target, double check the following:

- Did you specify the correct port world-wide name?
- If you are creating a virtual target from a subset of LUN(s) of a physical device, did you specify the correct fibre-channel (physical) LUN(s) and iSCSI (virtual) LUN(s)?
- If using an access list to control access to the virtual target, did you specify the correct initiator(s)? If you are not using an access list to restrict access, did you specify **all-initiator-permit** to insure all initiators have access?
- If restricting access to a particular interface(s), did you specify the correct gigabit Ethernet interface(s)?

## Useful show Commands for Debugging Static iSCSI Configuration

The output from the following commands reflects correctly established iSCSI sessions. Execute the same commands on your switch and compare with the output below to help identify possible issues:

**show iscsi session detail**

**show iscsi stats**

**show iscsi stats detail**

**show fcns data vsan 5**

**show flogi data vsan 5**

**show iscsi remote-node iscsi-session-detail tcp-parameters**

### show iscsi session detail

```
switch# show iscsi session detail
Initiator ign.1987-05.com.cisco.02.8cb3c18879bf356ce18e09679103235f.my-kayak (MY-KAYAK)
 Session #1 (index 84)
 Target ign.com.domainname.IPS-TEST.02-08.gw.2200002037c52d6d
 VSAN 5, ISID 00023d000054, TSID 135, Status active, no reservation
 Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
 MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
 DataSeqInOrder No, InitialR2T Yes, ImmediateData No
 Registered LUN 0, Mapped LUN 0
 Stats:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
 Number of connection: 1
 Connection #1
 Local IP address: 0xa011d64, Peer IP address: 0xa011d65
 CID 0, State: LOGGED_IN
 StatsSN 1356, ExpStatsSN 0
 MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
 CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
 AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
 Version Min: 0, Max: 0
 FC target: Up, Reorder PDU: No, Marker send: No (int 0)
 Received MaxRecvDSLen key: Yes

 Session #2 (index 85)
 Target ign.com.domainname.IPS-TEST.02-08.gw.2200002037c52e2e
 VSAN 5, ISID 00023d000055, TSID 135, Status active, no reservation
 Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
```

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

```
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
Stats:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
Number of connection: 1
Connection #1
 Local IP address: 0xa011d64, Peer IP address: 0xa011d65
 CID 0, State: LOGGED_IN
 StatSN 1356, ExpStatSN 0
 MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
 CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
 AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
 Version Min: 0, Max: 0
 FC target: Up, Reorder PDU: No, Marker send: No (int 0)
 Received MaxRecvDSLen key: Yes
```

```
Session #3 (index 86)
Target ign.com.domainname.IPS-TEST.02-08.gw.2200002037c52356
VSAN 5, ISID 00023d000056, TSID 135, Status active, no reservation
Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
Stats:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
Number of connection: 1
Connection #1
 Local IP address: 0xa011d64, Peer IP address: 0xa011d65
 CID 0, State: LOGGED_IN
 StatSN 1356, ExpStatSN 0
 MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
 CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
 AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
 Version Min: 0, Max: 0
 FC target: Up, Reorder PDU: No, Marker send: No (int 0)
 Received MaxRecvDSLen key: Yes
```

```
Session #4 (index 87)
Target ign.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
VSAN 5, ISID 00023d000057, TSID 135, Status active, no reservation
Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
Stats:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
Number of connection: 1
Connection #1
 Local IP address: 0xa011d64, Peer IP address: 0xa011d65
 CID 0, State: LOGGED_IN
 StatSN 1356, ExpStatSN 0
 MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
 CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
 AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
 Version Min: 0, Max: 0
 FC target: Up, Reorder PDU: No, Marker send: No (int 0)
 Received MaxRecvDSLen key: Yes
```

*Send comments to mdsfeedback-doc@cisco.com.*

## show iscsi stats

```
switch# sh iscsi stats iscsi2/7
iscsi2/7
 5 minutes input rate 3336 bits/sec, 417 bytes/sec, 0 frames/sec
 5 minutes output rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
 iSCSI statistics
 4112871 packets input, 4022464380 bytes
 303100 Command pdus, 3740086 Data-out pdus, 3815901300 Data-out bytes, 0
 fragments
 1283306 packets output, 778111088 bytes
 303069 Response pdus (with sense 3163), 195108 R2T pdus
 715480 Data-in pdus, 715214528 Data-in bytes
```

## show iscsi stats detail

```
switch# sh iscsi stats detail
iscsi2/7
 5 minutes input rate 3336 bits/sec, 417 bytes/sec, 0 frames/sec
 5 minutes output rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
 iSCSI statistics
 4113028 packets input, 4022586092 bytes
 303140 Command pdus, 3740200 Data-out pdus, 3816015476 Data-out bytes, 0
 fragments
 1283382 packets output, 778114736 bytes
 303109 Response pdus (with sense 3163), 195141 R2T pdus
 715480 Data-in pdus, 715214528 Data-in bytes
 iSCSI Forward:
 Command: 303140 PDUs (Received: 303140)
 Data-Out (Write): 3740200 PDUs (Received 3740200), 0 fragments, 3816015476 b
 ytes
 TMF Request: 0 (Received 28)
 FCP Forward:
 Xfer_rdy: 195141 (Received: 195141)
 Data-In: 715480 (Received: 715622), 715214528 bytes
 Response: 303109 (Received: 303322), with sense 3163
 TMF Resp: 0

 iSCSI Stats:
 Login: attempt: 16726, succeed: 114, fail: 16606, authen fail: 0
 Rcvd: NOP-Out: 36164, Sent: NOP-In: 36160
 NOP-In: 0, Sent: NOP-Out: 0
 TMF-REQ: 28, Sent: TMF-RESP: 0
 Text-REQ: 39, Sent: Text-RESP: 0
 SNACK: 0
 Unrecognized Opcode: 0, Bad header digest: 0
 Command in window but not next: 0, exceed wait queue limit: 0
 Received PDU in wrong phase: 0
 FCP Stats:
 Total: Sent: 4110679
 Received: 1281518 (Error: 0, Unknown: 0)
 Sent: PLOGI: 66367, Rcvd: PLOGI_ACC: 71, PLOGI_RJT: 66296
 PRLI: 71, Rcvd: PRLI_ACC: 71, PRLI_RJT: 0, Error resp: 0
 LOGO: 0, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
 ABTS: 87, Rcvd: ABTS_ACC: 0
 TMF REQ: 0
 Self orig command: 213, Rcvd: data: 142, resp: 213
 Rcvd: PLOGI: 614, Sent: PLOGI_ACC: 490
 LOGO: 197, Sent: LOGO_ACC: 111
 PRLI: 0, Sent: PRLI_ACC: 0
 ABTS: 183

 iSCSI Drop:
```

*Send comments to mdsfeedback-doc@cisco.com.*

```

Command: Target down 0, Task in progress 0, LUN map fail 0
 CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
 Persistent Resv 0 Data-Out: 0, TMF-Req: 0
FCP Drop:
 Xfer_rdy: 0, Data-In: 0, Response: 0

Buffer Stats:
 Buffer less than header size: 48475, Partial: 2524437, Split: 3550971
 Pullup give new buf: 48475, Out of contiguous buf: 0, Unaligned m_data: 0

```

## show fcns database

```

switch# sh fcns data vsan 5

VSAN 5:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x610002 N 20:0b:00:0b:be:77:72:42 scsi-fcp:init isc..w
0x6101e1 NL 22:00:00:20:37:c5:2d:6d (Seagate) scsi-fcp:target
0x6101e2 NL 22:00:00:20:37:c5:2e:2e (Seagate) scsi-fcp:target
0x6101e4 NL 22:00:00:20:37:c5:23:56 (Seagate) scsi-fcp:target
0x6101e8 NL 22:00:00:20:37:c5:26:0a (Seagate) scsi-fcp:target

Total number of entries = 5

```

## show flogi database

```

switch# sh flogi data vsan 5

INTERFACE VSAN FCID PORT NAME NODE NAME

fc1/12 5 0x6101e8 22:00:00:20:37:c5:26:0a 20:00:00:20:37:c5:26:0a
fc1/12 5 0x6101e4 22:00:00:20:37:c5:23:56 20:00:00:20:37:c5:23:56
fc1/12 5 0x6101e2 22:00:00:20:37:c5:2e:2e 20:00:00:20:37:c5:2e:2e
fc1/12 5 0x6101e1 22:00:00:20:37:c5:2d:6d 20:00:00:20:37:c5:2d:6d
iscsi2/8 5 0x610002 20:0b:00:0b:be:77:72:42 20:0a:00:0b:be:77:72:42

Total number of flogi = 5.

```

## show iscsi remote-node iscsi-session-detail tcp parameters

```

switch# sh iscsi remote-node iscsi-session-detail tcp-parameters
iSCSI Node name is ign.1987-05.com.cisco.02.8cb3c18879bf356ce18e09679103235f.my-kayak
 iSCSI alias name: MY-KAYAK
 Node WWN is 20:0a:00:0b:be:77:72:42 (dynamic)
 Member of vsans: 5
 Number of Virtual n_ports: 1

Virtual Port WWN is 20:0a:00:0b:be:77:72:42 (dynamic)
 Interface iSCSI 2/8, Portal group tag is 0x87
 VSAN ID 0, FCID 0x0
 No. of FC sessions: 1
 No. of iSCSI sessions: 1

iSCSI session details

Target node:
Statistics:
 PDU: Command: 0, Response: 0
 Bytes: TX: 0, RX: 0
 Number of connection: 1

```



*Send comments to mdsfeedback-doc@cisco.com.*

```

TCP parameters
 Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1026
 Path MTU 1500 bytes
 Current retransmission timeout is 310 ms
 Round trip time: Smoothed 179 ms, Variance: 33
 Advertized window: Current: 62 KB, Maximum: 62 KB, Scale: 0
 Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
 Congestion window: Current: 63 KB
VSAN ID 5, FCID 0x610002
No. of FC sessions: 4
No. of iSCSI sessions: 4

iSCSI session details

Target node: ign.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
 Number of connection: 1
TCP parameters
 Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
 Path MTU 1500 bytes
 Current retransmission timeout is 300 ms
 Round trip time: Smoothed 165 ms, Variance: 35
 Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
 Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
 Congestion window: Current: 63 KB

Target node: ign.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
 Number of connection: 1
TCP parameters
 Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
 Path MTU 1500 bytes
 Current retransmission timeout is 300 ms
 Round trip time: Smoothed 165 ms, Variance: 35
 Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
 Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
 Congestion window: Current: 63 KB

Target node: ign.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
 Number of connection: 1
TCP parameters
 Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
 Path MTU 1500 bytes
 Current retransmission timeout is 300 ms
 Round trip time: Smoothed 165 ms, Variance: 35
 Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
 Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
 Congestion window: Current: 63 KB

Target node: ign.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
 PDU: Command: 13, Response: 13
 Bytes: TX: 1344, RX: 0
 Number of connection: 1
TCP parameters
 Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
 Path MTU 1500 bytes

```

*Send comments to mdsfeedback-doc@cisco.com.*

```
Current retransmission timeout is 300 ms
Round trip time: Smoothed 165 ms, Variance: 35
Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
Congestion window: Current: 63 KB
```

## Fine Tuning/Troubleshooting IPS iSCSI TCP Performance

Generally there are two segments which will effect the iSCSI performance. First is the FC side flow control mechanism (Buffer to Buffer Credits, and the FC max frame size) Second is the TCP/IP side.

As in all TCP/IP-related throughput issues, the most important criteria are the Receive/Send Window Sizes on both TCP end points, RTT (Round Trip Time), actual available bandwidth between the TCP peers, the MSS (Maximum Segment Size) and the support for higher MTUs between the peers..

The following commands will give you information related to these criteria.

**show iscsi remote-node iscsi-session-detail tcp-parameters**

**show ips stats tcp interface gigabitethernet 2/1 detail**

**show interface iscsi <x/y>**

*show interface gigabitethernet <x/y>*

**show interface fc <x/y>**

**show iscsi remote-node fcp-session-detail**

The default MTU size of an ethernet network is 1500, while the FC networks generally support maximum frame sizes of 2148 bytes.. This means that an iSCSI gateway will need to chop the FC frames into two TCP segments or IP fragments while transferring form the FC side to the IP side depending on how this chopping is implemented within the device.

The IPS module adjusts the Receive Data Field Size that it advertises to its FC partner, according to the MTU that is configured on the corresponding Gigabit port of an iSCSI client.

If left to default MTU, the FC frame size from the Target device is decreased to match the maximum Ethernet frame size, so that the switching of the packet through the switch is swifter. Hence, one point of performance tuning is increasing the MTU of the IP network between the peers. In this setup there is one single Catalyst switch.

Jumbo support was enabled for the IPS ports, as well as the MTU for the VLAN corresponding to these ports was increased.

The second point is to increase the TCP window size of the iSCSI end points. Depending on the latency between the iSCSI client and IPS, this will need fine tuning. The switch's iSCSI configuration defines the TCP window size in kilobytes.

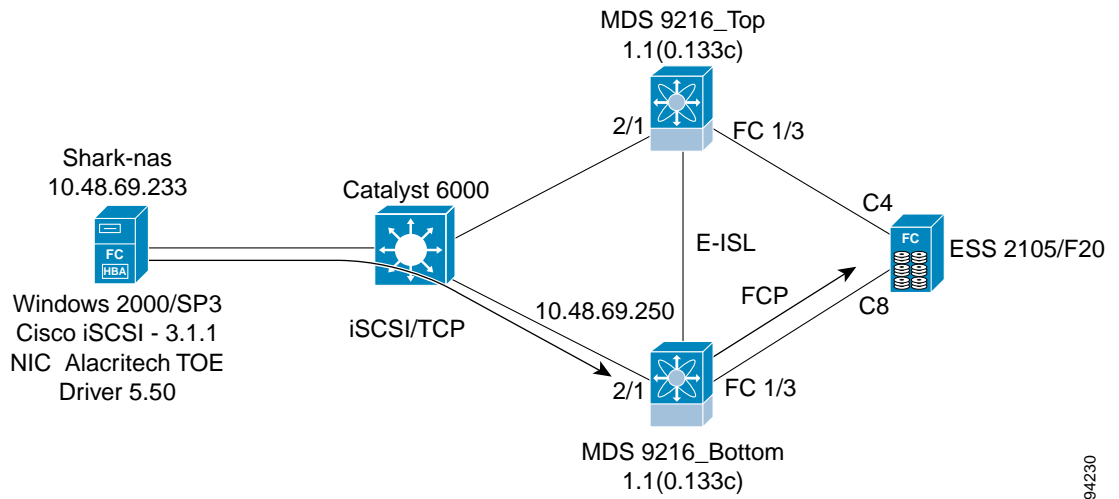
Any value starting with 64K (> 65535 = 0xFFFF bytes) will automatically trigger TCP window scaling according to RFC1323. The IPS TCP Window scaling begins only when the remote peer (iSCSI client in this case) requests it. This means that you need to configure the TCP stack of your client to trigger this functionality (see Figure 5-14).

For the FC side, depending on the direction of the traffic, the B2Bcredit of the ports corresponding to the input interfaces (feeding/receiving traffic to/from the iSCSI side) could be increased, especially in the case of local Gigabit Ethernet attached iSCSI clients.

Each of the above-mentioned commands are taken from a scenario in Figure 5-14. The important sections of the displays are highlighted/italicized or bolded.

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

**Figure 5-14 IPS Window Scaling**



## Lab Setup

This is the lab setup that was used in collecting the performance-related information.

The server was an IBM PentiumIII Server: Dual CPU @ 1.13 Ghz

The *tcp window-size* at both ends was set to 1MB (1024K).

The IBM ESS Shark had a hardcoded B2B value of 64 (not configurable).

The *fcrxbbcredit* on the corresponding switch port (fc1/3) was set to the same value.

The C4 and C8 represented the corresponding port WWPNS for the IBM Shark storage subsystem. See below for full WWPNS:

C4 → 50:05:07:63:00:c4:94:4c (in VSAN 778)

C8 → 50:05:07:63:00:c8:94:4c (in VSAN 777)

## Configuration from the Bottom MDS

The following example is the configuration for the 9216 switch shown in Figure 5-14.

```
iscsi initiator name ign.1987-05.com.cisco:02.75af2f95624c.shark-nas
pWWN 20:05:00:0c:30:6c:24:42
 vsan 777
 vsan 778

iscsi virtual-target name shark_nas
pWWN 50:05:07:63:00:c8:94:4c fc-lun 0000 iscsi-lun 0000 secondary-pwwn
50:05:07:63:00:c4:94:4c
pWWN 50:05:07:63:00:c8:94:4c fc-lun 0001 iscsi-lun 0001 secondary-pwwn
50:05:07:63:00:c4:94:4c
initiator ign.1987-05.com.cisco:02.75af2f95624c.shark-nas permit

interface GigabitEthernet2/1
ip address 10.48.69.251 255.255.255.192
iscsi authentication none
```

*Send comments to mdsfeedback-doc@cisco.com.*

```
no shutdown
vrrp 1
priority 110
address 10.48.69.250
(This is the iSCSI target IP address for the Windows iSCSI client.)

no shutdown

interface iscsi2/1
tcp pmtu-enable
tcp window-size 1024
(To increase the receive window size of the IPS module (in kilobytes).)

tcp sack-enable
no shutdown
```

To verify the connectivity between your client and the IPS iSCSI service:

```
MDS_BOTTOM# show ips stats tcp interface gig 2/1
TCP Statistics for port GigabitEthernet2/1
Connection Stats
 0 active openings, 24 accepts
 0 failed attempts, 0 reset received, 24 established
Segment stats
 7047380 received, 56080130 sent, 0 retransmitted
 0 bad segments received, 0 reset sent

TCP Active Connections
 Local Address Remote Address State Send-Q Recv-Q
 10.48.69.250:3260 10.48.69.233:1026 ESTABLISH 0 0
 10.48.69.250:3260 10.48.69.233:1057 ESTABLISH 34560 0
 0.0.0.0:3260 0.0.0.0:0 LISTEN 0 0

MDS_BOTTOM# show flogi database vsan 777

INTERFACE VSAN FCID PORT NAME NODE NAME

fc1/3 777 0x610000 50:05:07:63:00:c8:94:4c 50:05:07:63:00:c0:94:4c
iscsi2/1 777 0x610001 20:05:00:0c:30:6c:24:42 20:00:00:0c:30:57:5e:c2

Total number of flogi = 2.

MDS_BOTTOM# show fcns dabase vsan 777

VSAN 777:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x610000 N 50:05:07:63:00:c8:94:4c (IBM) scsi-fcp:target fc..
0x610001 N 20:05:00:0c:30:6c:24:42 scsi-fcp:init isc..w

Total number of entries = 2
MDS_BOTTOM#
MDS_BOTTOM# show module
Mod Ports Module-Type Model Status
--- ---
1 16 1/2 Gbps FC/Supervisor DS-X9216-K9-SUP active *
2 8 IP Storage Module DS-X9308-SMIP ok

Mod Sw Hw World-Wide-Name(s) (WWN)
--- ---

```

*Send comments to mdsfeedback-doc@cisco.com.*

```
1 1.1(0.133c) 1.0 20:01:00:0c:30:57:5e:c0 to 20:10:00:0c:30:57:5e:c0
2 1.1(0.133c) 0.2 20:41:00:0c:30:57:5e:c0 to 20:48:00:0c:30:57:5e:c0
```

| Mod | MAC-Address(es)                        | Serial-Num  |
|-----|----------------------------------------|-------------|
| 1   | 00-0b-be-f8-7f-00 to 00-0b-be-f8-7f-04 | JAB070804Q3 |
| 2   | 00-05-30-00-a8-56 to 00-05-30-00-a8-62 | JAB070205am |

\* this terminal session

MDS\_BOTTOM#

MDS\_BOTTOM# **show iscsi remote**

iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas

iSCSI alias name: SHARK-NAS

Node WWN is 20:00:00:0c:30:57:5e:c2 (dynamic)

Member of vsans: 777, 778

Number of Virtual n\_ports: 1

Virtual Port WWN is 20:05:00:0c:30:6c:24:42 (configured)

Interface iSCSI 2/1, Portal group tag: 0x1001

VSAN ID 778, FCID 0x7c0000

VSAN ID 777, FCID 0x610001

MDS\_BOTTOM# **show iscsi local**

target: shark\_nas

Port WWN 50:05:07:63:00:c8:94:4c

(This is the port of the Shark connected to mds\_bottom.)

Secondary PWWN 50:05:07:63:00:c4:94:4c

(This is the port of the Shark connected to mds\_top.)

Configured node

No. of LU mapping: 2

iSCSI LUN: 0000, FC LUN: 0000

iSCSI LUN: 0001, FC LUN: 0001

No. of initiators permitted: 1

initiator iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas is permitted

all initiator permit is disabled

MDS\_BOTTOM#

MDS\_BOTTOM# **show interface iscsi 2/1**

iscsi2/1 is up

Hardware is GigabitEthernet

Port WWN is 20:41:00:0c:30:57:5e:c0

Admin port mode is iSCSI

Port mode is iSCSI

Speed is 1 Gbps

Number of iSCSI session: 2, Number of TCP connection: 2

Configured TCP parameters

Local Port is 3260

PMTU discover is **enabled** (default)

(This is especially required if there may be devices without jumbo support in the path. The initial TCP 3-way handshake will establish a session with a high MSS value (provided both the IPS module and the iSCSI client are configured/capable) even if there are devices without jumbo frame support in the path. Without PMTU discovery, this will create problems.)

Keepalive-timeout 60

Initial-retransmit-time 300

(If there is high delay between the peers, this is one of the parameters that can be adjusted. There's no real formula, rather use trial and error to find the optimum value for your network. Try lower values as well as higher ones, and get hints from the **show ips stats tcp** display.)

Max-retransmissions 8

Window-size **1024000**

Sack is enabled

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

```
Forwarding mode: pass-thru
5 minutes input rate 410824 bits/sec, 51353 bytes/sec, 1069 frames/sec
5 minutes output rate 581291520 bits/sec, 72661440 bytes/sec, 53302 frames/sec
iSCSI statistics
 1072393 packets input, 51482588 bytes
 1072305 Command pdus, 0 Data-out pdus, 0 Data-out bytes, 0 fragments
 53430805 packets output, 72837086312 bytes
 1072273 Response pdus (with sense 9), 0 R2T pdus
 52358444 Data-in pdus, 70272402880 Data-in bytes
```

```
MDS_BOTTOM# show iscsi remote initiator ign.1987-05.com.cisco:02.75af2f95624c.shark-nas
iscsi tcp
iSCSI Node name is ign.1987-05.com.cisco:02.75af2f95624c.shark-nas
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:57:5e:c2 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1
```

```
Virtual Port WWN is 20:00:00:0c:30:57:5e:c2 (configured)
Interface iSCSI 2/1, Portal group tag is 0x1001
 VSAN ID 0, FCID 0x 0
 No. of FC sessions: 1
 No. of iSCSI sessions: 1
```

#### iSCSI session details

```
Target node:
Statistics:
 PDU: Command: 0, Response: 0
 Bytes: TX: 0, RX: 0
 Number of connection: 1
TCP parameters
 Local 10.48.69.250:3260, Remote 10.48.69.233:1026
 Path MTU: 1500 bytes
 Retransmission timeout: 300 ms
 Round trip time: Smoothed 150 ms, Variance: 31
 Advertized window: Current: 998 KB, Maximum: 1000 KB, Scale: 4
 Peer receive window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4
 Congestion window: Current: 12 KB
VSAN ID 777, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1
```

#### iSCSI session details

```
Target node: shark_nas
Statistics:
 PDU: Command: 392051, Response: 392042
 Bytes: TX: 25692593152, RX: 0
 Number of connection: 1
TCP parameters
 Local 10.48.69.250:3260, Remote 10.48.69.233:1057
 Path MTU: 1500 bytes
 Retransmission timeout: 300 ms
 Round trip time: Smoothed 2 ms, Variance: 1
```

(Watch out for these numbers. The above output is for a TCP session that goes only through one Gigabit Ethernet switch. When there are multiple router hops, as well as WAN links in the middle, the RTT will grow, and the variance will fluctuate with higher values. You may need to adjust the Retransmission timeout.)

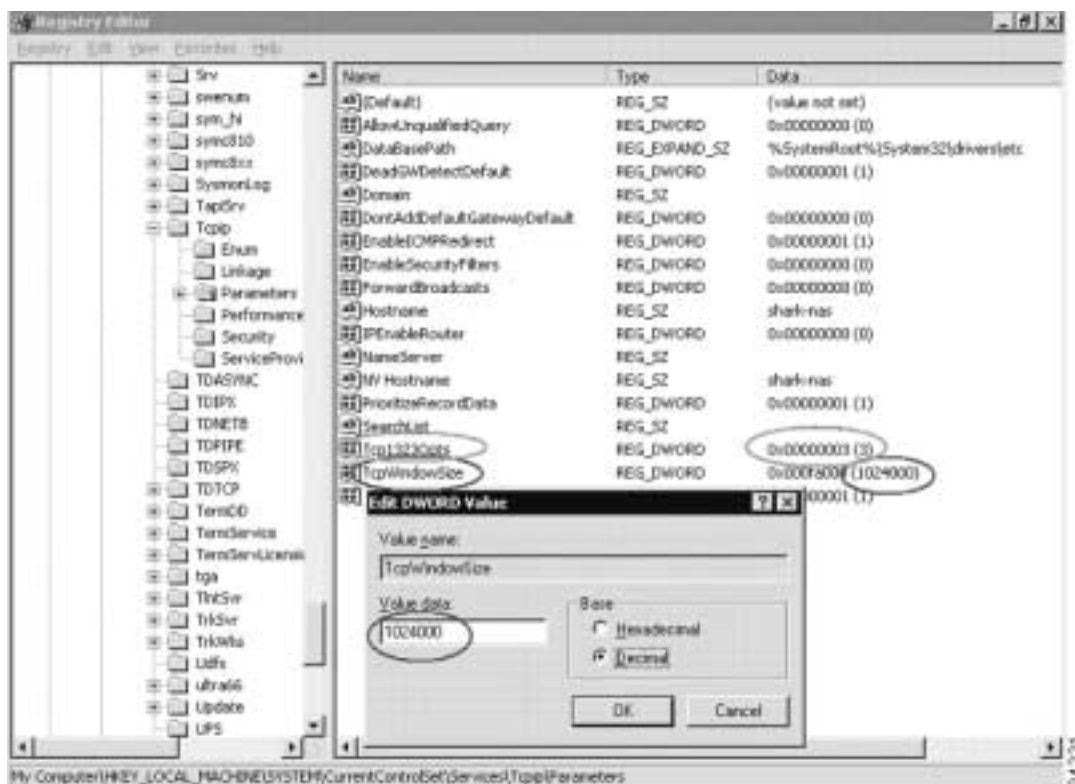
Advertized window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4  
(This is the window size set on the Windows Client. See Figure 5-15.)

Peer receive window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4

*Send comments to mdsfeedback-doc@cisco.com.*

(This is the window size set on the IPS iscsi interface. See Figure 5-15.)

**Figure 5-15 Congestion window: Current: 24 KB**



## Changing TCP Parameters

To change TCP parameters in the Windows registry use the Registry Parameters above as an example. Setting the `Tcp1323Opts` (marked green) to 3, sets two bits ON, one for Window Scaling, the other for the Timestamp Option. We are only interested in the Window Scaling here. User should be aware that editing registry is a high risk operation and can render the System unusable requiring a reinstall of the whole operating system !! Only advanced users should do this.

```
MDS_BOTTOM# show interface gigabitethernet 2/1
GigabitEthernet2/1 is up
 Hardware is GigabitEthernet, address is 0005.3000.a85a
 Internet address is 10.48.69.251/26
 MTU 1500 bytes, BW 1000000 Kbit
```

(Better throughput can be achieved if the MTU of both the client NIC, as well as the IPS Gigabit interface is changed for higher MTU, provided the network in the middle supports jumbo frames.)

```
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
```

*Send comments to mdsfeedback-doc@cisco.com.*

```
5 minutes input rate 3957384 bits/sec, 494673 bytes/sec, 6716 frames/sec
5 minutes output rate 609420144 bits/sec, 76177518 bytes/sec, 53267 frames/sec
6979248 packets input, 514206826 bytes
 0 multicast frames, 0 compressed
 0 input errors, 0 frame, 0 overrun 0 fifo
55551272 packets output, 79456286344 bytes, 0 underruns
 0 output errors, 0 collisions, 0 fifo
 0 carrier errors
```

```
MDS_BOTTOM# show interface fc 1/3
```

```
fc1/3 is up
```

```
Hardware is Fibre Channel
Port WWN is 20:03:00:0c:30:57:5e:c0
Admin port mode is auto, trunk mode is on
Port mode is F, FCID is 0x610000
Port vsan is 777
Speed is 1 Gbps
Transmit B2B Credit is 64
```

(This depends on the storage device; it can not be changed.)

```
Receive B2B Credit is 64
```

(This is the switch's receive; the default is 16 for F/FL ports.)

```
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 524947584 bits/sec, 65618448 bytes/sec, 49382 frames/sec
5 minutes output rate 470432 bits/sec, 58804 bytes/sec, 988 frames/sec
64560099 frames input, 85630621884 bytes
 0 discards, 3 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
1291861 frames output, 76739928 bytes
 0 discards, 0 errors
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
3 output OLS, 3 LRR, 0 NOS, 2 loop inits
```

```
MDS_BOTTOM# show ips stats tcp interface gigabit 2/1 detail
```

```
TCP Statistics for port GigabitEthernet2/1
```

```
TCP send stats
56252632 segments, 76746280484 bytes
56100434 data, 152173 ack only packets
1 control (SYN/FIN/RST), 0 probes, 24 window updates
0 segments retransmitted, 0 bytes
```

(The lower the better. Increasing values would show that the IP network in the middle has issues, or that the TCP peer has problems ACK'ing the data that IPS sends to it.)

```
0 retransmitted while on ethernet send queue, 0 packets split
```

(Packets Split shows the IP level fragmentation; would increase if the MTU of this interface is higher than the MSS of the iSCSI client; for example, client MTU default 1500 => MSS=1460, but IPS Gigabit MTU changed to 2500).)

```
3 delayed acks sent
TCP receive stats
7068115 segments, 1061853 data packets in sequence, 54245464 bytes in sequence
0 predicted ack, 187 predicted data
0 bad checksum, 0 multi/broadcast, 0 bad offset
0 no memory drops, 0 short segments
0 duplicate bytes, 0 duplicate packets
0 partial duplicate bytes, 0 partial duplicate packets
0 out-of-order bytes, 0 out-of-order packets
0 packet after window, 0 bytes after window
0 packets after close
7067879 acks, 76746255713 ack bytes, 0 ack toomuch, 21 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
```



*Send comments to mdsfeedback-doc@cisco.com.*

```

5980106 window updates, 0 window probe
50 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
TCP Connection Stats
 0 attempts, 24 accepts, 24 established
 22 closed, 2 drops, 0 conn drops
 0 drop in retransmit timeout, 0 drop in keepalive timeout
 0 drop in persist drops, 0 connections drained
TCP Miscellaneous Stats
 7054414 segments timed, 7067879 rtt updated
 0 retransmit timeout, 0 persist timeout
 19 keepalive timeout, 19 keepalive probes
TCP SACK Stats
 0 recovery episodes, 54218621 data packets, 77791012992 data bytes
 0 data packets retransmitted, 0 data bytes retransmitted
 1 connections closed, 0 retransmit timeouts
TCP SYN Cache Stats
 24 entries, 24 connections completed, 0 entries timed out
 0 dropped due to overflow, 0 dropped due to RST
 0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
 0 abort due to no memory, 0 duplicate SYN, 2 no-route SYN drop
 0 hash collisions, 0 retransmitted

TCP Active Connections
 Local Address Remote Address State Send-Q Recv-Q
 10.48.69.250:3260 10.48.69.233:1026 ESTABLISH 0 0
 10.48.69.250:3260 10.48.69.233:1057 ESTABLISH 29296 0
 0.0.0.0:3260 0.0.0.0:0 LISTEN 0 0

MDS_BOTOM# show iscsi remote-node fcp-session-detail
iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:6c:24:42 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1

Virtual Port WWN is 20:00:00:0c:30:6c:24:42 (configured)
Interface iSCSI 2/1, Portal group tag is 0x1001
VSAN ID 0, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1

FCP Session details

Target FCID: 0x000000 (S_ID of this session: 0x000000)
 pWWN: 00:00:00:00:00:00:00:00
 nWWN: 00:00:00:00:00:00:00:00
Session state: INIT
1 iSCSI sessions share this FC session
Target:
Negotiated parameters
 RcvDataFieldSize 2048 our_RcvDataFieldSize 1392
Will be set to maximum 2048 if the MTU is increased on the Gigabit interface
corresponding to this iscsi remote-node. See below for an example.
 MaxBurstSize 0, EMPD: FALSE
 Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
 PDU: Command: 0, Response: 0
VSAN ID 777, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1

FCP Session details

Target FCID: 0x610000 (S_ID of this session: 0x610001)

```

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

```
pWWN: 50:05:07:63:00:c8:94:4c
Verify that the local port, rather than a remote that is reached via an ISL link
is used for the storage target by the above field to avoid suboptimal access to storage.
You can see that C8 is the locally attached port of the shark storage subsystem.
```

```
nWWN: 50:05:07:63:00:c8:94:4c
Session state: LOGGED_IN
1 iSCSI sessions share this FC session
Target: shark_nas
Negotiated parameters
RcvDataFieldSize 2048 our_RcvDataFieldSize 1392
MaxBurstSize 0, EMPD: FALSE
Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
PDU: Command: 0, Response: 1612007
MDS_BOTTOM#
```

The following example shows the effect of changing the Gigabit MTU on FC RcvDataFieldSize.

```
interface GigabitEthernet2/1
ip address 10.48.69.249 255.255.255.192
iscsi authentication none
switchport mtu 2440
no shutdown
vrrp 1
address 10.48.69.250
no shutdown

MDS_Top# show iscsi remote-node iscsi-session-detail tcp-parameters
iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:6c:24:42 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1

Virtual Port WWN is 20:00:00:0c:30:6c:24:42 (configured)
Interface iSCSI 2/1, Portal group tag is 0x1001
VSAN ID 0, FCID 0x 0
No. of FC sessions: 1
No. of iSCSI sessions: 1

iSCSI session details

Target node:
Statistics:
PDU: Command: 0, Response: 0
Bytes: TX: 0, RX: 0
Number of connection: 1
TCP parameters
Local 10.48.69.250:3260, Remote 10.48.69.233:1026
Path MTU: 2440 bytes
Retransmission timeout: 420 ms
Round trip time: Smoothed 94 ms, Variance: 83
Advertized window: Current: 999 KB, Maximum: 1000 KB, Scale: 4
Peer receive window: Current: 1024 KB, Maximum: 1024 KB, Scale: 4
Congestion window: Current: 11 KB
VSAN ID 777, FCID 0x700003
No. of FC sessions: 1
No. of iSCSI sessions: 1

iSCSI session details

Target node: shark_nas
```

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

```

Statistics:
 PDU: Command: 11, Response: 11
 Bytes: TX: 2152, RX: 0
 Number of connection: 1
TCP parameters
 Local 10.48.69.250:3260, Remote 10.48.69.233:1040
 Path MTU: 2440 bytes
 Retransmission timeout: 370 ms
 Round trip time: Smoothed 47 ms, Variance: 81
 Advertized window: Current: 999 KB, Maximum: 1000 KB, Scale: 4
 Peer receive window: Current: 1024 KB, Maximum: 1024 KB, Scale: 4
 Congestion window: Current: 12 KB

MDS_Top# show iscsi remote-node fcp-session-detail
iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
 iSCSI alias name: SHARK-NAS
 Node WWN is 20:00:00:0c:30:6c:24:42 (dynamic)
 Member of vsans: 777, 778
 Number of Virtual n_ports: 1

Virtual Port WWN is 20:00:00:0c:30:6c:24:42 (configured)
Interface iSCSI 2/1, Portal group tag is 0x1001
 VSAN ID 0, FCID 0x 0
 No. of FC sessions: 1
 No. of iSCSI sessions: 1

FCP Session details

Target FCID: 0x000000 (S_ID of this session: 0x000000)
 pWWN: 00:00:00:00:00:00:00:00
 nWWN: 00:00:00:00:00:00:00:00
Session state: INIT
1 iSCSI sessions share this FC session
Target:
Negotiated parameters
 RcvDataFieldSize 2048 our_RcvDataFieldSize 2048
 MaxBurstSize 0, EMPD: FALSE
 Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
 PDU: Command: 0, Response: 0
VSAN ID 777, FCID 0x700003
 No. of FC sessions: 1
 No. of iSCSI sessions: 1

FCP Session details

Target FCID: 0x700000 (S_ID of this session: 0x700003)
 pWWN: 50:05:07:63:00:c4:94:4c
 nWWN: 50:05:07:63:00:c4:94:4c
Session state: LOGGED_IN
1 iSCSI sessions share this FC session
Target: shark_nas
Negotiated parameters
 RcvDataFieldSize 2048 our_RcvDataFieldSize 2048
 MaxBurstSize 0, EMPD: FALSE
 Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
 PDU: Command: 0, Response: 11

```

To get the real benefit of this increased MTU and higher FC Frame Size, the path between the iSCSI client and the IPS iSCSI interface (as well as the host NIC) has to be capable of supporting this high MTU.

*Send comments to mdsfeedback-doc@cisco.com.*

If you do not have access to the host, one way to see if the host is also configured for high MTU/MSS (as well as the path in the middle) is to check the split packets field in the **show ips stats tcp** display :

However this is a generic display for all TCP sessions. That is, if you have some Hosts with high MTU-capable NICs, and some others without, it may be difficult to assess which is which.

```
MDS_Top# show ips stats tcp interface gig 2/1 detail (truncated output)
TCP Statistics for port GigabitEthernet2/1
 TCP send stats
 10 segments, 240 bytes
 5 data, 5 ack only packets
 0 control (SYN/FIN/RST), 0 probes, 0 window updates
 0 segments retransmitted, 0 bytes
 0 retransmitted while on ethernet send queue, 0 packets split
.....

 TCP Active Connections
 Local Address Remote Address State Send-Q Recv-Q
 10.48.69.250:3260 10.48.69.233:1026 ESTABLISH 0 0
 10.48.69.250:3260 10.48.69.233:1040 ESTABLISH 0 0
 0.0.0.0:3260 0.0.0.0:0 LISTEN 0 0
```

MDS\_Top#

Afterward, traffic starts flowing from the FC storage towards the server that is connected via iSCSI to the IPS.

```
MDS_Top# show ips stats tcp interface gig 2/1 detail (truncated output)
TCP Statistics for port GigabitEthernet2/1
 TCP send stats
 715535 segments, 943511612 bytes
 712704 data, 2831 ack only packets
 0 control (SYN/FIN/RST), 0 probes, 0 window updates
 0 segments retransmitted, 0 bytes
 0 retransmitted while on ethernet send queue, 345477 packets split
.....
```



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

INDEX

---

## C

### CLI

launching from Fabric View 1-11

### configuration guide

related documents viii

### connectivity

verifying 1-10

---

## D

documentation feedback ix

---

## E

### end-to-end connectivity

See connectivity

---

## F

fabric configuration, analyzing 1-10

---

## L

### launching

CLI 1-11

---

## M

merging zones 1-11

---

## O

ordering documentation ix

---

## S

switch health 1-9

---

## T

### TAC

Escalation Center x, xi

web site x

### Technical Assistance Center

See TAC x

traceroute 1-11

### troubleshooting

with traceroute 1-11

---

## V

### verifying

fabric configuration 1-10

zone configuration 1-11

---

## Z

### zones

merging 1-11