Cisco Systems

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco MDS 9000 Family Troubleshooting Guide, Release 2.x

Cisco MDS SAN-OS for Release 2.0(1b) through Release 2.1(2b)

December 2005

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

Text Part Number: OL-9076-01

Cisco MDS 9000 Family Troubleshooting Guide, Release 2.x

CONTENTS

New and Changed Information

This chapter provides release-specific information for each new and changed troubleshooting guideline for the Cisco MDS SAN-OS Release 2.x software. The *Cisco MDS 9000 Family Troubleshooting Guide, Release 2.x* is updated to address each new and changed guideline in the Cisco MDS SAN-OS Release 2.x software. The latest version of this document is available at the following Cisco Systems website: http://www.cisco.com/en/US/products/ps5989/prod_troubleshooting_guides_list.html

Tip The troubleshooting guides created for previous releases are also listed in the website mentioned above. Each guide addresses the features introduced in or available in those releases. Select and view the troublehsooting guide pertinent to the software installed in your switch.

To check for additional information about Cisco MDS SAN-OS Release 2.x, refer to the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website: http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_release_notes_list.html

Table 1 summarizes the new and changed features for the *Cisco MDS 9000 Family Troubleshooting Guide, Release 2.x*, and tells you where they are documented. The table includes a brief description of each new feature and the release in which the change occurred.

Note This updated version of the *Cisco MDS 9000 Family Troubleshooting Guide, Release 2.x* has been reorganized from earlier versions to better address the most common troubleshooting issues in Cisco SAN-OS Release 2.x.

Table 1 New and Changed Features for Release 2.x

Feature	Description	Changed in Release	Where Documented
Upgrades/Downgrades	Added troubleshooting Cisco SAN-OS upgrades, downgrades and reboots and bootflash recovery.	All releases	Chapter 2, "Troubleshooting Installs, Upgrades, and Reboots"
Hardware	Added troubleshooting Fans, power supplies and clock modules.	All releases	Chapter 3, "Troubleshooting Hardware"
Licenses	Added troubleshooting license issues.	1.3(1)	Chapter 4, "Troubleshooting Licensing"

Table 1 New and Changed Features for Release 2.x (continued)

Feature	Description	Changed in Release	Where Documented
Domains, FSPF	Added troubleshooting options for domains and FSPF.	All releases	Chapter 6, "Troubleshooting Ports"
Inter-VSAN Routing (IVR) Enhancements	Describes troubleshooting IVR, including IVR NAT and IVR auto topology.	2.1(2b)	Chapter 8, "Troubleshooting IVR"

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 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 Installs, Upgrades, and Reboots	Describes how to identify and resolve problems that might occur when installing, upgrading, or rebooting Cisco MDS 9000 Family hardware.
Chapter 3	Troubleshooting Hardware	Describes how to identify and resolve problems that might occur when replacing modules, fans, chassis, power supplies or other hardware.
Chapter 4	Troubleshooting Cisco Fabric Services	Describes procedures used to troubleshoot Cisco Fabric Services (CFS) problems.
Chapter 5	Troubleshooting Licensing	Describes procedures used to troubleshoot licensing issues.
Chapter 6	Troubleshooting Ports	Describes how to identify and resolve problems that might occur when using port interfaces.
Chapter 7	Troubleshooting VSANs, Domains, and FSPF	Describes how to identify and resolve problems that might occur when using Virtual Storage Area Networks (VSANs).
Chapter 8	Troubleshooting IVR	Describes how to debug and resolve Inter-VSAN Routing (IVR) configuration issues.
Chapter 9	Troubleshooting Zones and Zone Sets	Describes how to identify and resolve problems that might occur while implementing zones and zone sets.

Chapter	Title	Description
Chapter 10	Troubleshooting IP Storage Services	Describes how to identify and resolve problems that might occur when using IP Services.
Chapter 11	Troubleshooting IPsec	Describes procedures used to troubleshoot IP security (IPsec) and Internet Key Exchange (IKE) encryption issues.
Chapter 12	Troubleshooting Fabric Manager Problems	Describes procedures used to troubleshoot Fabric Manager.
Appendix A	Before Contacting Technical Support	Describes the steps to perform before calling for technical support with any Cisco MDS 9000 Family product.
Appendix B	Troubleshooting Tools and Methodology	Describes the troubleshooting tools and methodology available for the Cisco MDS 9000 Family product.

Document Conventions

Command descriptions use these conventions:

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

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS SAN-OS Documentation Locator at:
http://www.cisco.com/en/US/products/ps5989/products_documentation_roadmap09186a00804500c1.html.
For information on IBM TotalStorage SAN Volume Controller Storage Software for the Cisco MDS 9000 Family, refer to the IBM TotalStorage Support website:
http://www.ibm.com/storage/support/2062-2300/

Release Notes

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

Compatibility Information

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

Regulatory Compliance and Safety Information

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

Hardware Installation

- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9020 Fabric Switch Hardware Installation Guide*

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco Fabric Manager

- *Cisco MDS 9000 Family Fabric Manager Quick Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Fabric Manager Online Help*

Command-Line Interface

- *Cisco MDS 9000 Family Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family CLI Quick Configuration Guide*
- *Cisco MDS 9000 Family CLI Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Quick Command Reference*
- *Cisco MDS 9020 Fabric Switch Configuration Guide and Command Reference*
- *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide*

Troubleshooting and Reference

- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family MIB Quick Reference*
- *Cisco MDS 9020 Fabric Switch MIB Quick Reference*
- *Cisco MDS 9000 Family CIM Programming Reference*
- *Cisco MDS 9000 Family System Messages Reference*
- *Cisco MDS 9020 Fabric Switch System Messages Reference*

Installation and Configuration Note

- *Cisco MDS 9000 Family SSM Configuration Note*
- *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Note*

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Send documentation comments to mdsfeedback-doc@cisco.com

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/techsupport

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

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

Product Documentation DVD

Cisco documentation and additional literature are available in the Product Documentation DVD package, which may have shipped with your product. The Product Documentation DVD is updated regularly and may be more current than printed documentation.

The Product Documentation DVD is a comprehensive library of technical product documentation on portable media. The DVD enables you to access multiple versions of hardware and software installation, configuration, and command guides for Cisco products and to view technical documentation in HTML. With the DVD, you have access to the same documentation that is found on the Cisco website without being connected to the Internet. Certain products also have .pdf versions of the documentation available.

The Product Documentation DVD is available as a single unit or as a subscription. Registered Cisco.com users (Cisco direct customers) can order a Product Documentation DVD (product number DOC-DOCDVD=) from Cisco Marketplace at this URL:

http://www.cisco.com/go/marketplace/

Ordering Documentation

Beginning June 30, 2005, registered Cisco.com users may order Cisco documentation at the Product Documentation Store in the Cisco Marketplace at this URL:

http://www.cisco.com/go/marketplace/

Nonregistered Cisco.com users can order technical documentation from 8:00 a.m. to 5:00 p.m. (0800 to 1700) PDT by calling 1 866 463-3487 in the United States and Canada, or elsewhere by calling 011 408 519-5055. You can also order documentation by e-mail at tech-doc-store-mkpl@external.cisco.com or by fax at 1 408 519-5001 in the United States and Canada, or elsewhere at 011 408 519-5001.

Documentation Feedback

You can rate and provide feedback about Cisco technical documents by completing the online feedback form that appears with the technical documents on Cisco.com.

You can send comments about Cisco documentation to mdsfeedback-doc@cisco.com.

Send documentation comments to mdsfeedback-doc@cisco.com

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

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

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

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

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

http://www.cisco.com/go/psirt

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

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

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com

  An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- Nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Tip We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.*x* through 8.*x*.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

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

The link on this page has the current PGP key ID in use.

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support & Documentation Website

The Cisco Technical Support & Documentation website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, at this URL:

http://www.cisco.com/techsupport

Access to all tools on the Cisco Technical Support & Documentation website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

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

Note Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support & Documentation website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)
EMEA: +32 2 704 55 55
USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

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

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

  http://www.cisco.com/packet

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

  http://www.cisco.com/go/iqmagazine

  or view the digital edition at this URL:

  http://ciscoiq.texterity.com/ciscoiq/sample/

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

  http://www.cisco.com/ipj

- Networking products offered by Cisco Systems, as well as customer support services, can be obtained at this URL:

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

- Networking Professionals Connection is an interactive website for networking professionals to share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

  http://www.cisco.com/discuss/networking

- World-class networking training is available from Cisco. You can view current offerings at this URL:

  http://www.cisco.com/en/US/learning/index.html

CHAPTER 1

Troubleshooting Overview

This chapter introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that may occur when configuring and using the Cisco MDS 9000 Family of multilayer directors and fabric switches.

This chapter includes the following sections:

Overview of the Troubleshooting Process

To troubleshoot your fabric environment, follow these general steps:

Step 1 Gather information that defines the specific symptoms.

Step 2 Identify all potential problems that could be causing the symptoms.

Step 3 Systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear.

To identify the possible problems, you need to use a variety of tools and understand the overall storage environment. For this reason, this guide describes a number of general troubleshooting tools in Appendix B, "Troubleshooting Tools and Methodology," including those that are specific to the Cisco MDS 9000 Family. This chapter also provides a plan for investigating storage issues. See other chapters in this book for detailed explanations of specific issues.

Overview of Best Practices

Best practices are the recommended steps you should take to ensure the proper operation of your fabric. Each chapter includes a section on best practices for the covered Cisco SAN-OS features. We recommend the following general best practices for most SAN fabrics:

- Maintain a consistent Cisco SAN-OS release across all your Cisco MDS switches.
- Refer to the release notes for your Cisco SAN-OS release for the latest features, limitations, and caveats.
- Enable system message logging. See the "Overview of Symptoms" section on page 1-8.
- Troubleshoot any new configuration changes after implementing the change.
- Use Fabric Manager and Device Manager to proactively manage your fabric and detect possible problems before they become critical.

Troubleshooting Basics

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 access some, but not all, of the logical unit numbers (LUNs) on an existing subsystem, then fabric-specific issues (such as FSPF, ISLs, or FCNS) do not need to be investigated. The fabric components can therefore be eliminated from possible causes of the problem.

This section contains the following topics:

Troubleshooting Guidelines

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

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.

Answer the following questions to determine the status of your installation:

- 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?

To discover a SAN problem, use the following general SAN troubleshooting steps:

Step 1 Gather information on problems in your fabric. See the "Gathering Information Using Common Fabric Manager Tools and CLI Commands" section on page 1-3.

Step 2 Verify physical connectivity between your switches and end devices. See the "Verifying Basic Connectivity" section on page 1-4.

Step 3 Verify registration to your fabric for all SAN elements. See the "Verifying SAN Element Registration" section on page 1-5.

Step 4 Verify the configuration for your end devices (storage subsystems and servers).

Step 5 Verify end-to-end connectivity and fabric configuration. See the "Fibre Channel End-to-End Connectivity" section on page 1-5.

Gathering Information Using Common Fabric Manager Tools and CLI Commands

This section highlights the Fabric Manager tools and CLI commands that are commonly used to troubleshoot problems within your fabric. These tools and commands are a subset of what you may use to troubleshoot your specific problem. Each chapter may include tools and commands specific to the symptoms and possible problems.

Common Fabric Manager Tools

Use the following navigation paths in Fabric Manager or Device Manager to access common troubleshooting information:

- Overview of switch status—In Fabric Manager, click the **Switch Health Analysis** icon.
- End-to-end connectivity—In Fabric Manager, click the **End-to-End Connectivity Analysis** icon.
- Fabric configuration— In Fabric Manager, click the **Fabric Configuration Analysis** icon.
- Module status—In Device Manager, choose **Physical > Modules**.
- Cisco SAN-OS version—In Device Manager, choose **Physical > System.**
- View logs—In Device Manager, choose **Logs > FM Server** or **Logs > Switch Resident**.
- View Fabric Manager events—In Fabric Manager, click the **Events** tab in the map pane.
- Interface status—In Fabric Manager, choose **Switches > Interfaces** and select the port type you are interested in.
- View name server information— In Device Manager, choose **FC > Name Server**.
- View FLOGI information—In Fabric Manager, choose **Switches > Interfaces > FC Physical > FLOGI**.
- Analyze the results of merging zones – In Fabric Manager, choose **Zone > Merge Analysis**.

Fabric Manager and Device Manager also provide the following tools to proactively monitor your fabric:

- ISL performance—In Fabric Manager, click the **ISL Performance** icon.
- Network monitoring—In Device Manage, click the **Summary** tab.
- Performance monitoring—In Fabric Manager, choose **Performance > Start Collection**.

Common CLI Commands

Issue the following commands and examine the outputs:

- **show module**
- **show version**
- **show running-config**
- **show logging log**
- **show interfaces brief**
- **show fcns**
- **show flogi**
- **show hardware internal errors**
- **show zoneset active**
- **show accounting log**

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

Verifying Basic Connectivity

Answer the following questions to verify basic connectivity between your end devices:

- Are you using the correct fiber (SM or MM)?
- Did you check for a broken fiber?
- Is the Fibre Channel port LED on the connected module green, and do the LEDs on any host bus adapter (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 a 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 an N port?

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 to 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.

In Fabric Manager, choose **Tools > Ping** or **Tools > Traceroute** (or use the **fcping** or **fctrace** CLI commands) to verify connectivity. See the "FC Ping and FC Traceroute" section on page B-4.

Verifying SAN Element Registration

Answer the following questions to verify that your end devices are registered to the fabric:

- Are the HBAs and subsystem ports successfully registered with the fabric name server?
  - In Device Manager, choose **FC > Name Server**.
  - In the CLI, use the **show fcns** commands.
- Does the correct pWWN for the HBAs and the storage subsystem ports show up on the correct port in the FLOGI database?
  - In Fabric Manager, choose **Switches > Interfaces > FC Physical > FLOGI**.
  - In the CLI, use the **show flogi** commands.
- Are the HBA and storage subsystem on the same VSAN?
  - In Fabric Manager, choose **End Devices** and verify the VSAN IDs are identical.
  - From the CLI, use the **show vsan membership** command.
- Does any single zone contain both devices?
  - In Fabric Manager, choose the **Zone > Edit Full Zone Database** and select the active zone set (in bold) for the VSAN that contains the end devices. Verify that both devices are members of the same zone.
  - From the CLI, use the **show zoneset active** command.

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 port WWN (pWWN) or FC ID in its logs?
- Does the subsystem list the host's pWWN or FC ID 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 N port?

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

Fabric Issues

Answering the following questions will help to determine the status of the fabric configuration:

- Are both the 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 or FC IDs.
- Is the zone correctly configured and part of the active configuration or zone set 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?

Port Issues

Initial tasks to perform while investigating port connectivity issues include:

- Verify correct media: copper or optical; single-mode (SM) or multimode (MM).
- Is the media broken or damaged?
- Is the LED on the switch green?
- Is the active LED on the HBA for the connected device on?

Basic port monitoring using Device Manager begins with the visual display in the Device View. (See Figure 1-1.) Port display descriptions include:

- Green box: A successful fabric login has occurred; the connection is active.
- Red X: A small form-factor pluggable (SFP) transceiver is present but there is no connection. This could indicate a disconnected or faulty cable, or no active device connection.
- Red box: An SFP is present but fabric login (FLOGI) has failed. Typically there is a mismatch in port or fabric parameters with the neighboring device. For example, a port parameter mismatch would occur if a node device were connected to a port configured as an E port. An example of a fabric parameter mismatch would be differing timeout values.
- Yellow box: In Device Manager, a port has been selected.
- Gray box: The port is administratively disabled.
- Black box: An SFP is not present.

Figure 1-1 Device Manager: Device View

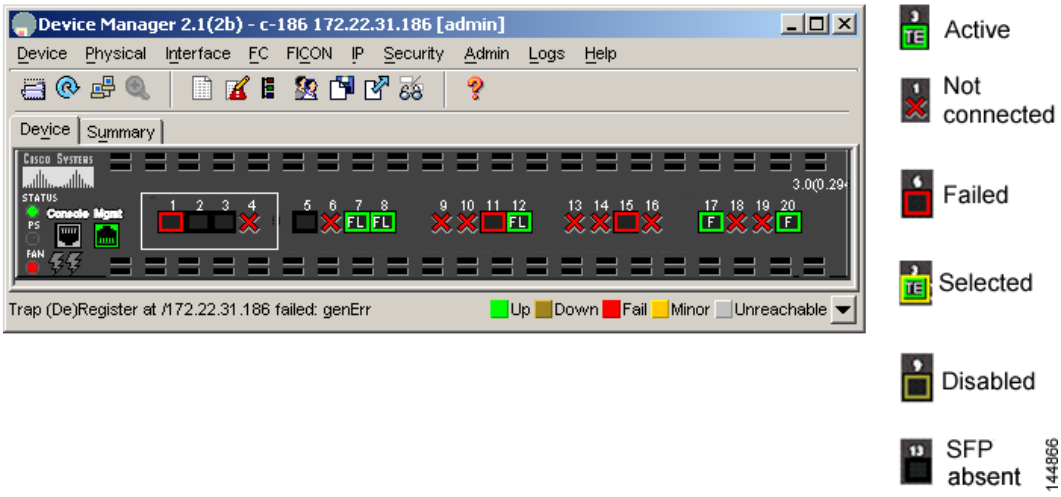

Device Manager: Summary View

In Device Manager, selecting the Summary View expands the information available for port monitoring. (See Figure 1-2.) The display includes:

- VSAN assignment
- For N ports, the port World Wide Name (pWWN) and Fibre Channel ID (FC ID) of the connected device
- For ISLs, the IP address of the connected switch
- Speed
- Frames transmitted and received
- Percentage utilization for the CPU, dynamic memory, and Flash memory

Figure 1-2 Device Manager: Summary View

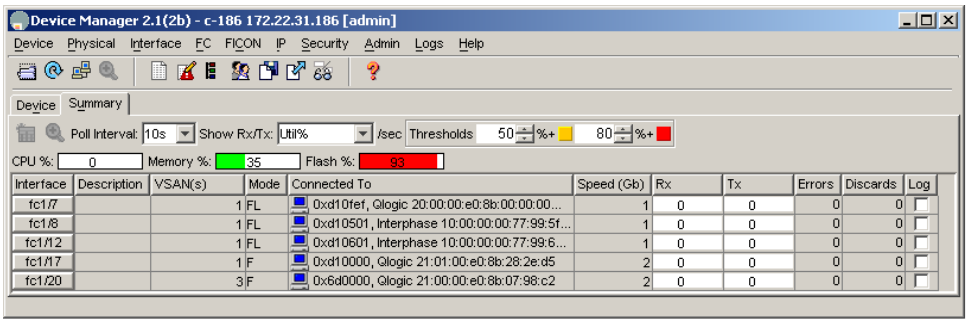

Device Manager: Port Selection

To drill down for additional port information, use the Device View or Summary View. Select and double-click any port. The initial display shows administrative settings for Mode, Speed, and Status, plus current operational status, failure cause, and date of the last configuration change.

Additional tabs include:

- Rx BB Credit—Configure and view buffer-to-buffer credits (BB_credits).
- Other—View PortChannel ID, WWN, and maximum transmission unit (MTU), and configure maximum receive buffer size.
- FLOGI—View FC ID, pWWN, nWWN, BB_credits, and class of service for N port connections.
- ELP—View pWWN, nWWN, BB_credits, and supported classes of service for ISLs.
- Trunk Config—View and configure trunk mode and allowed VSANs.
- Trunk Failure—View the failure cause for ISLs.
- Physical—Configure beaconing; view SFP information.
- Capability—View current port capability for hold-down timers, BB credits, maximum receive buffer size.

Primary Troubleshooting Flowchart

The flowchart in Figure 1-3 shows the overall troubleshooting process. Begin any troubleshooting investigation by checking one of the following four areas:

- Physical port issues
- Physical switch issues
- Fx port issues
- Fabric services

Figure 1-3 Troubleshooting Process Flowchart

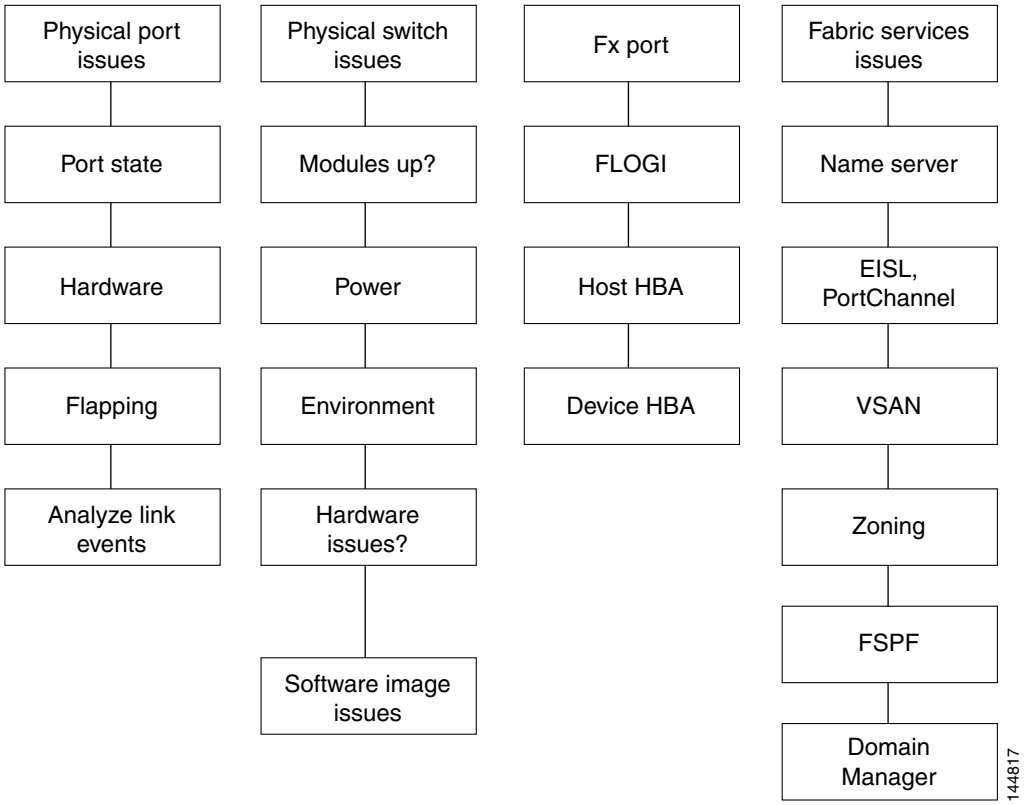

Overview of Symptoms

The symptom-based troubleshooting approach provides multiple ways to diagnose and resolve problems. By using multiple entry points with links to solutions, this guide best serves users who may have identical problems that are perceived by different indicators. Search this guide in PDF form, use the index, or rely on the symptoms and diagnostics listed in each chapter as entry points to access necessary information in an efficient manner.

Using a given a set of observable symptoms on a Fibre Channel SAN, it is important to be able to diagnose and correct software configuration issues and inoperable hardware components, so that the problems are resolved with minimal disruption to the SAN environment. Those problems and corrective actions include:

- Identify key Cisco MDS troubleshooting tools.
- Obtain and analyze Fibre Channel protocol traces using RSPAN on the CLI.
- Identify or rule out physical port issues.
- Identify or rule out switch module issues.
- Diagnose and correct Fx port issues.
- Diagnose and correct issues on the data path.
- Diagnose and correct advanced services issues.
- Recover from switch upgrade failures.
- Diagnose and resolve Fabric Manager and Device Manager configuration problems.
- Obtain core dumps and other diagnostic data for use by the TAC.

System Messages

The system software sends these syslog (system) messages to the console (and, optionally, to a logging server on another system) during operation. Not all messages indicate a problem with your system. Some messages are purely informational, while others might help diagnose problems with links, internal hardware, or the system software.

This section contains the following topics:

System Message Text

Message-text is a text string that describes the condition. This portion of the message might contain detailed information about the event, including terminal port numbers, network addresses, or addresses that correspond to locations in the system memory address space. Because the information in these variable fields changes from message to message, it is represented here by short strings enclosed in square brackets ([ ]). A decimal number, for example, is represented as [dec].

```
PORT-3-IF_UNSUPPORTED_TRANSCEIVER: Transceiver for interface [chars] is not supported.
```

Use this string to find the matching system message in the *Cisco MDS 9000 Family System Messages Reference.*

Each system message is followed by an explanation and recommended action. The action may be as simple as "No action required." It may involve a fix or a recommendation to contact technical support as shown in the following example:

Error Message `PORT-3-IF_UNSUPPORTED_TRANSCEIVER: Transceiver for interface [chars] is not supported.`

Explanation Transceiver (SFP) is not from an authorized vendor.

Recommended Action Enter the **show interface transceiver** CLI command or similar Fabric Manager/Device Manager command to determine the transceiver being used. Please contact your customer support representative for a list of authorized transceiver vendors.

Syslog Server Implementation

The syslog facility allows the Cisco MDS 9000 Family platform to send a copy of the message log to a host for more permanent storage. This can be useful if the logs need to be examined over a long period of time or when the Cisco MDS switch is not accessible.

This example will demonstrate how to configure a Cisco MDS switch to utilize the syslog facility on a Solaris platform. Although a Solaris host is being used, syslog configuration on all UNIX and Linux systems is very similar.

Syslog uses the concept of a facility to determine how it should be handled on the syslog server (the Solaris system in this example), and the message severity. Therefore, different message severities can be handled differently by the syslog server. They could be logged to different files or e-mailed to a particular user. Specifying a severity determines that all messages of that level and greater severity (lower number) will be acted upon.

Note The Cisco MDS messages should be logged to a different file from the standard syslog file so that they cannot be confused with other non-Cisco syslog messages. The logfile should not be located on the / file system, to prevent log messages from filling up the / file system.
Syslog Client: switch1
Syslog Server: 172.22.36.211 (Solaris)
Syslog facility: local1
Syslog severity: notifications (level 5, the default)
File to log MDS messages to: /var/adm/MDS_logs

Implementing Syslog with Fabric Manager

To configure system message logging servers, follow these steps:

Step 1 In Fabric Manager, choose **Switches > Events > Syslog** and click the **Servers** tab in the Information pane.

In Device Manager, choose **Logs > Syslog > Setup** and click the **Servers** tab in the Syslog dialog box.

Step 2 Click **Create Row** in Fabric Manager or **Create** in Device Manager to add a new syslog server.

Step 3 Enter the name or IP address in dotted decimal notation (for example, 192.168.2.12) of the syslog server in the Name or IP Address field.

Step 4 Set the message severity threshold by clicking the **MsgSeverity** radio button and set the facility by clicking the **Facility** radio button.

Step 5 Click **Apply Changes** in Fabric Manager or click **Create** in Device Manager to save and apply your changes.

Step 6 If CFS is enabled in Fabric Manager for the syslog feature, click **CFS** and commit these changes to propagate the configuration through the fabric.

Device Manager allows you to view event logs on your local PC as well as those on the switch. For a permanent record of all events that occur on the switch, you should store these messages off the switch. To do this the Cisco MDS switch must be configured to send syslog messages to your local PC and a syslog server must be running on that PC to receive those messages. These messages can be categorized into four classes:

- Hardware—Line card or power supply problems
- Link incidents—FICON port condition changes
- Accounting—User change events
- Events—All other events

Note You should avoid using PCs that have IP addresses randomly assigned to them by DHCP. The switch continues to use the old IP address unless you manually change it; however the Device Manager prompts you if it does detect this situation. UNIX workstations have a built-in syslog server. You must have root access (or run the Cisco syslog server as setuid to root) to stop the built-in syslog daemon and start the Cisco syslog server.

Implementing Syslog with the CLI

To configure a syslog server using the CLI, follow these steps:

Step 1 Configure the Cisco MDS switch:

```
switch1# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# logging server 172.22.36.211 6 facility local1
```

To display the configuration:

```
switch1# show logging server
Logging server: enabled
{172.22.36.211}
     server severity: notifications
     server facility: local1
```

Step 2 Configure the syslog server:

a. Modify /etc/syslog.conf to handle local1 messages. For Solaris, there needs to be at least one tab between the facility.severity and the action (/var/adm/MDS_logs).

```
#Below is for the MDS 9000 logging
local1.notice /var/adm/MDS_logs
```

b. Create the log file.

```
#touch /var/adm/MDS_logs
```

c. Restart syslog.

```
# /etc/init.d/syslog stop
# /etc/init.d/syslog start
syslog service starting.
```

d. Verify syslog started.

```
# ps -ef |grep syslogd
     root 23508 1 0 11:01:41 ? 0:00 /usr/sbin/syslogd
```

Step 3 Test the syslog server by creating an event on the Cisco MDS switch . In this case, port fc1/2 was bounced and the following was listed on the syslog server. Notice that the IP address of the switch is listed in brackets.

```
# tail -f /var/adm/MDS_logs
Sep 17 11:07:41 [172.22.36.142.2.2] : 2004 Sep 17 11:17:29 pacific:
%PORT-5-IF_DOWN_INITIALIZING: %$VSAN 1%$ Interface fc1/2 is down (Initializing)
Sep 17 11:07:49 [172.22.36.142.2.2] : 2004 Sep 17 11:17:36 pacific: %PORT-5-IF_UP:
%$VSAN 1%$ Interface fc1/2 is up in mode TE
Sep 17 11:07:51 [172.22.36.142.2.2] : 2004 Sep 17 11:17:39 pacific:
%VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(dhcp-171-71-49-125.cisco.com)
```

Troubleshooting with Logs

Cisco SAN-OS generates many types of system messages on the switch and sends them to a syslog server. These messages can be viewed using Fabric Manager or the CLI to determine what events may have led up to the current problem condition you are facing.

This section contains the following topics:

Viewing Logs with Fabric Manager

Fabric Manager and Device Manager present concise views of the generated system messages and other logged events:

- In Device Manager, click **Logs** to set up and view logs.
- In Fabric Manager, select the **Logs** tab at the bottom of the map pane to view log information.
- Learn to use Threshold Manager to alert you that critical statistics have exceeded a set threshold.

Viewing Logs with the CLI

The following CLI commands are available to access and view logs on a switch:

```
Musky-9506# show logging ?

console Show console logging configuration
info Show logging configuration
last Show last few lines of logfile
level Show facility logging configuration
logfile Show contents of logfile
module Show module logging configuration
monitor Show monitor logging configuration
nvram Show NVRAM log
server Show server logging configuration
<cr> Carriage Return
```

Example 1-1 shows an example of the **show logging** CLI command output.

Example 1-1 show logging Command

```
Musky-9506# show logging server
Logging server: enabled
{10.91.51.204}
server severity: critical
server facility: user
```

Viewing the Log from the Supervisor

You can view system messages from Device Manager if Device Manager is running from the same workstation as the Fabric Manager Server. Choose **Logs > Events > current** to view the system messages on Device Manager. The columns in the events table are sortable. In addition, you can use the Find button to locate text within the table.

You can view switch-resident logs even if you have not set up your local syslog server or your local PC is not in the switch's syslog server list. Because of memory constraints, these logs will wrap when they reach a certain size. The switch syslog has two logs: an NVRAM log that holds a limited number of critical and greater messages and a nonpersistent log that contains notice or greater severity messages. Hardware messages are part of these logs.

Use the **show logging** CLI command to view the logs on the supervisor.

Viewing NVRAM logs

System messages that are priority 0, 1, or 2 are logged into NVRAM on the supervisor module. After a switch reboots, you can display these syslog messages in NVRAM using the show logging nvram CLI command. See Example 1-2.

Example 1-2 Show logging nvram

```
switch# show logging nvram
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-PS_OK: Power supply 2 ok (Serial
number )
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-PS_FANOK: Fan in Power supply 2 ok
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-FANMOD_FAN_OK: Fan module 1 (Front fan) ok
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-FANMOD_FAN_OK: Fan module 2 (Rear fan) ok
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-CHASSIS_CLKMODOK: Chassis clock module A ok
```

```
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-CHASSIS_CLKMODOK: Chassis clock module B ok
2005 Sep 16 13:19:20 172.20.150.82 %PLATFORM-2-CHASSIS_CLKSRC: Current chassis clock
source is clock-A
2005 Sep 16 13:19:36 172.20.150.82 %PLATFORM-2-PFM_STDBY_BIOS_STUCK: standby supervisor
bios failed
2005 Sep 16 13:20:19 172.20.150.82 %IMAGE_DNLD-SLOT13-2-IMG_DNLD_STARTED:  Module image
download process. Please wait until completion...
2005 Sep 16 13:20:32 172.20.150.82 %IMAGE_DNLD-SLOT13-IMG_DNLD_COMPLETE:  Module image
download process. Download successful.
2005 Sep 16 15:44:46 172.20.150.82 %PLATFORM-2-PFM_STDBY_BIOS_STUCK: standby supervisor
bios failed
2005 Sep 16 15:44:53 172.20.150.82 %PLATFORM-2-MOD_ALL_PWRDN_NOXBAR: All modules powered
down due to non-availability of xbar modules
2005 Sep 16 15:45:41 172.20.150.82 %PLATFORM-2-MOD_PWRUP_XBAR: Modules powered up due to
xbar availability
2005 Sep 18 15:12:07 172.20.150.82 %MODULE-2-MOD_FAIL: Initialization of module 14
(serial: JAB092501FC) failed
```

Contacting Customer Support

If you are unable to solve a problem after using the troubleshooting suggestions in this guide, contact a customer service representative for assistance and further instructions. Before you call, have the following information ready to help your service provider assist you as quickly as possible:

- Date you received the switch
- Chassis serial number (located on a label on the right side of the rear panel of the chassis)
- Type of software and release number
- Maintenance agreement or warranty information
- Brief description of the problem
- Brief explanation of the steps you have already taken to isolate and resolve the problem

After you have collected this information, see the "Obtaining Technical Assistance" section on page xxv.

For more information on steps to take before calling Technical Support, see the "Before Contacting Technical Support" section on page A-1.

CHAPTER 2

Troubleshooting Installs, Upgrades, and Reboots

This chapter describes how to identify and resolve problems that might occur when installing, upgrading, or restarting Cisco MDS 9000 Family products. It includes the following sections:

Overview

Each Cisco MDS 9000 switch ships with an operating system (Cisco SAN-OS) that consists of two images—the kickstart image and the system image. There is also a module image if the Storage Services Module (SSM) is present.

Installations, upgrades, and reboots are ongoing parts of SAN maintenance activities. It is important to minimize the risk of disrupting ongoing operations when performing these operations in production environments, and to know how to recover quickly when something does go wrong.

Note For documentation purposes, we use the term upgrade in this document. However, upgrade refers to both upgrading and downgrading your switch, depending on your needs.

Best Practices

This sections lists the best practices for Cisco SAN-OS software installations, image upgrade and downgrade procedures, and reboots and includes the following topics:

Best Practices for Installations

Follow these best practices guidelines for installing Cisco SAN-OS software images:

- Server availability—Ensure that an FTP or TFTP server is available.
- Compatibility check from CLI—Use the **show install all impact** CLI command to verify that the new image is healthy and the impact that new load will have on any hardware with regards to compatibility. Check for compatibility.
- Compatibility check using Device Manager—Choose **Admin > Show Image Version** in the Device Manager to view information on images in the directories of the MDS file system.

Best Practices for Upgrading

Not all images need to be updated during an upgrade. Use the following checklist to prepare for an upgrade:

Checklist	Checkoff
Copy the new Cisco SAN-OS image onto your supervisor modules in bootflash: or slot0:.	❑
Save your running configuration to the startup configuration.	❑
Backup a copy of your configuration to a remote TFTP server.	❑
Schedule your upgrade during an appropriate maintenance window for your fabric.	❑

After you have completed the checklist, you are ready to upgrade the switches in your fabric.

Note It is normal for the active supervisor to become the standby supervisor during an upgrade.

Follow these best practices guidelines for upgrading and downgrading Cisco SAN-OS software images:

- Read the Cisco SAN-OS Release Notes for the release you are upgrading or downgrading to. Cisco SAN-OS Release Notes are available at the following website:

  http://cisco.com/en/US/products/ps5989/prod_release_notes_list.html
- Ensure that an FTP or TFTP server is available.

- Copy the startup-config to a snapshot config in NVRAM. This creates a backup copy of the startup-config.
  - In Device Manager, Choose Admin > Copy Configuration and select the **startupConfig** radio button for the From: field and the **serverFile** radio button for the To: field. Set the other fields and click **Apply**.
- From the CLI, use the **copy nvram:startup-config nvram-snapshot-config** CLI command.
- Where possible, choose to do a nondisruptive upgrade. You can nondisruptively upgrade to Cisco SAN-OS Release 2.x from any Cisco SAN-OS software release beginning with Release 1.3(x). If you are running an older version of Cisco SAN-OS, upgrade to Release 1.3(x) and then Release 2.x.
- Establish a PC serial connection to each supervisor console to record upgrade activity to a file. This catches any error messages or problems during bootup.
- In Fabric Manager, choose **Tools > Other > Software Install** or click the **Software Install** icon on the toolbar to use the Software Install Wizard.
- From the CLI, use the **install all** [{**asm-sfn** | **kickstart** | **ssi** | **system**} *URL*] command to run a complete script, test the images, and verify the compatibility with the hardware. See the "Installing Cisco SAN-OS Software from the CLI" section on page 2-10. Using the **install all** command offers the following advantages:
  - You can upgrade the entire switch using the least disruptive procedure with just one command.
  - You can receive descriptive information on the intended changes to your system before you continue with the command.
  - You have the option to cancel the command. Once the effects of the command are presented, you can continue or cancel when you see this question (the default is no):

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

  - You can view the progress of this command on the console, Telnet, and SSH screens.
  - The image integrity is automatically checked, including the running kickstart and system images.
  - The command performs a platform validity check to verify that a wrong image is not used. For example, the command verifies that an MDS 9500 Series image is not used inadvertently to upgrade an MDS 9200 Series switch.
  - After issuing the **install all** command, if any step in the sequence fails, the command completes the step in progress and ends.

    For example, if a switching module fails to be updated for any reason (for example, due to an unstable fabric state), then the command sequence disruptively updates that module and ends. In such cases, you can verify the problem on the affected switching module and upgrade the other switching modules.

Best Practices for Reboots

There are three different types of system restarts:

- Recoverable—A process restarts and service is not affected.
- Unrecoverable—A process has restarted more than the maximum 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 the system.

Schedule the reboot to avoid possible disruption of services during critical business hours.

Note Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM. You can view this log at any time with the **show logging nvram** CLI command.

Disruptive Module Upgrades

Software upgrades for the SSM, MPS-14/2 or the IP Storage (IPS) services modules are disruptive. These modules use a rolling upgrade install mechanism where the modules are upgraded in sequence. After the first module upgrade finishes, and before the next module upgrade begins, Cisco SAN-OS introduces a time delay to ensure that all applications in the module reach a steady state. The IPS modules require a five-minute delay before the next IPS module upgrade can guarantee a stable state.

SSM supports nondisruptive upgrades for the Layer 1 and Layer 2 protocols under the following conditions:

- SSM is running Cisco SAN-OS Release 2.1(2) or later and upgrading to a later release.
- The SSM hardware has the ELPD image for Release 2.1(2) installed. Use the s**how version module** <*module number*> **epld** CLI command and verify that the epld version is 0x07 or later.
- You have turned off all Layer 3 services on the SSM by deprovisioning the DPPs for Layer 3 service.

Troubleshooting Fabric Manager Installations

This section describes possible problems and solutions for a Fabric Manager installation failure. Fabric Manager requires the appropriate version Sun JAVA JRE installed, based on the Fabric Manager release. Table 2-1 shows the recommended JRE for Fabric Manager 2.x releases.

Table 2-1 Fabric Manager and Recommended JRE Version

Fabric Manager Release	Recommended JRE Version
2.0(1b) through 2.1(1b)	1.4.2_05
2.1(2) or later	1.5.0

Fabric Manager and Device Manager do not operate properly with JRE 1.4.2_03 on Windows 2003.

Symptom Fabric Manager or Device Manager will not start.

Table 2-2 Fabric Manager or Device Manager Will Not Start

Symptom	Possible Cause	Solution
Device Manager will not start.	Device Manager proxied through Fabric Manager Server.	Uncheck the **Proxy SNMP through FM Server** check box in the Device Manager startup dialog box and restart Device Manager.
Fabric Manager will not start.	Using incorrect Fabric Manager Server.	Verify that you are choosing the appropriate Fabric Manager Server from the FMServer pull-down menu. If you have not already done so, download Fabric Manager Server.
	Fabric Manager Server not running.	On a Windows PC, click **Start > Control Panel > Administrative Tools > Services** to verify that Fabric Manager Server and Fabric Manager database have started. The default setting for the Fabric Manager Server is that the server is automatically started when the PC is rebooted.
	Incompatible JRE version.	Verify that you have the correct JRE version installed for the Fabric Manager release you installed. Refer to the release notes for the software version you installed to determine which JRE version is compatible.
	Improperly installed.	If the problem remains, then remove the application using the Cisco MDS 9000/Uninstall program, then reinstall Fabric Manager.

Verifying Cisco SAN-OS Software Installations

In Fabric Manager, you can watch the progress of your software installation using the Software Install Wizard. From the CLI, you can use use the **show install all status** command to watch the progress of your software installation.

You can also use the **show install all status** CLI command to view the on-going **install all** command or the log of the last installed **install all** command from a console, SSH, or Telnet session.

This command presents the **install all** output on both the active and standby supervisor module even if you are not connected to the console terminal. It only displays the status of an **install all** command that is issued from the CLI (not the GUI). See Example 2-1.

Example 2-1 install all Command Output

```
switch# show install all status
There is an on-going installation... <---------------------- in progress installation
Enter Ctrl-C to go back to the prompt.
Verifying image bootflash:/b-1.3.0.104
-- SUCCESS
Verifying image bootflash:/i-1.3.0.104
-- SUCCESS
Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS
Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
```

```
-- SUCCESS
Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS
switch# show install all status
This is the log of last installation. <----------------- log of last install
Verifying image bootflash:/b-1.3.0.104
-- SUCCESS
Verifying image bootflash:/i-1.3.0.104
-- SUCCESS
Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS
Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
-- SUCCESS
Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS
```

Troubleshooting Cisco SAN-OS Software Upgrades and Downgrades

This section discusses possible causes and solutions for a software installation upgrade or downgrade failure. It includes the following symptoms:

Software Installation Reports an Incompatibility

Symptom The software installation reports an incompatibility.

Table 2-3 Software Installation Report Incompatibility

Symptom	Possible Cause	Solution
The software installation reports an incompatibility.	The running image may have a feature enabled that is not compatible with the proposed new image.	Review the incompatibility issues displayed by either the Fabric Manager Software Install Wizard or the **install all** CLI command. Correct any problems and retry the installation. See the "Diagnosing Compatibility Issues" section on page 2-6.
		Verify what features are enabled on your switch and disable any features that may not be compatible with your new image. Refer to the appropriate release notes for both images.

Diagnosing Compatibility Issues

To view the results of a dynamic compatibility check, use the **show incompatibility system** ***bootflash:filename*** CLI command.

Use the **show incompatibility** CLI command for diagnosis when the **install all** CLI command warns of compatibility issues.

During an attempted upgrade, the **install all** CLI command may return the following warning:

```
Warning: The startup config contains commands not supported by the system image; as a
result, some resources might become unavailable after an install.
Do you wish to continue? (y/ n) [y]: n
```

Use the **show incompatibility** CLI command to identify the problem.

Message 1 indicates that the remote SPAN (RSPAN) feature is in use, but it is not supported by the image that was installed. The incompatibility is strict because continuing the upgrade might cause the switch to move into an inconsistent state—that is, configured features might stop working.

```
switch# show incompatibility system bootflash:running-image
The following configurations on active are incompatible with the system image
1) Feature Index : 67 , Capability : CAP_FEATURE_SPAN_FC_TUNNEL_CFG
Description : SPAN - Remote SPAN feature using fc-tunnels
Capability requirement : STRICT
```

Message 2 indicates that the Fibre Channel tunnel feature is not supported in the new image. The RSPAN feature uses Fibre Channel tunnels.

```
2) Feature Index : 119 , Capability : CAP_FEATURE_FC_TUNNEL_CFG
Description : fc-tunnel is enabled
Capability requirement : STRICT
```

Software Installation Ends with Error

Symptom The software installation ends with an error.

Table 2-4 Software Installation Ends with Error

Problem	Possible Cause	Solution
The installation ends with an error.	The standby supervisor module bootflash: file system does not have sufficient space to accept the updated image.	Remove unnecessary files from the filesystem. In Device Manager, choose **Admin > Flash Files** and delete unnecessary files. From the CLI, use the **delete** command.
	The specified system and kickstart images are not compatible.	Check the output of the installation process for details on the incompatibility. Possibly update the kickstart image before updating the system image.
	The **install all** command is issued on the standby supervisor module.	Issue the command on the active supervisor module only.
	A module was inserted while the upgrade was in progress.	Restart the installation. See the "Installing SAN-OS Software Using Fabric Manager" section on page 2-9 or the "Installing Cisco SAN-OS Software from the CLI" section on page 2-10.
	The fabric or switch was configured while the upgrade was in progress.	Wait until the upgrade is complete before configuring the switch. In Device Manager, choose **Admin > CFS** or from the CLI, use the **show cfs lock** command to check that there are no CFS commit operations in progress.
	The switch experienced a power disruption while the upgrade was in progress.	Restart the installation. See the "Installing SAN-OS Software Using Fabric Manager" section on page 2-9 or the "Installing Cisco SAN-OS Software from the CLI" section on page 2-10.
	Incorrect software image path specified.	Specify the entire path for the remote location accurately.
	Another installation is already in progress.	Verify the state of the switch at every stage and restart the installation after 10 seconds. If you restart the installation within the 10-second span, the command is rejected with an error message indicating that an installation is currently in progress.

Installing SAN-OS Software Using Fabric Manager

To use the Software Install Wizard to install a new software image using Fabric Manager, follow these steps:

Step 1 Open the Software Install Wizard by clicking its icon in the toolbar (see Figure 2-1).

Figure 2-1 Software Install Wizard Icon

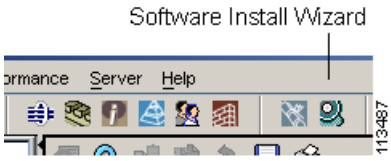

You see the Software Install Wizard.

Step 2 Select the switches you want to install images on. You must select at least one switch in order to proceed. Click **Next**.

Step 3 Optionally, check the **Skip Image Download** check box and click **Next** to use images that are already downloaded (the file is already on the bootflash: file system). Proceed to Step 7.

Step 4 Click the row under the System, Kickstart, Asm-sfn, or ssi columns to enter image URIs.You must specify at least one image for each switch to proceed.

Step 5 Check the active (and standby, if applicable) bootflash: file system on each switch to see if there is enough space for the new images. You can see this information in the Flash Space column.

This screen shows the active (and standby, if applicable) bootflash: memory space on each switch, and shows the status (whether there is enough space for the new images). If any switch has insufficient space, you cannot proceed. Deselect the switch without enough bootflash: memory by going back to the first screen and unchecking the check box for that switch.

Step 6 Click **Next**. You see the Select Download Image screen.

Step 7 Double-click the table cell under System, Kickstart, Asm-sfn, or Ssi and select from a drop-down list of images available in the bootflash: file system on each switch. You must select at least one image for each switch to proceed.

Note There is no limit on the number of switches you can upgrade. However, the upgrade is a serial process; that is, only a single switch is upgraded at a time.

Step 8 Click **Next**. You see the final verification screen.

Step 9 Click **Finish** to start the installation or click **Cancel** to leave the installation wizard without installing new images.

Note On hosts where the TFTP server cannot be started, a warning is displayed. The TFTP server may not start because an existing TFTP server is running or because access to the TFTP port 69 has been denied for security reasons (the default setting on LINUX). In these cases, you cannot transfer files from the local host to the switch.

Note Before exiting the session, be sure the upgrade process is complete. The wizard will display a status as it goes along. Check the lower left-hand corner of the wizard for the status message `Upgrade Finished`. First, the wizard displays the message `Success` followed a few seconds later by `InProgress Polling`. Then the wizard displays a second message `Success` before displaying the final `Upgrade Finished`.

Installing Cisco SAN-OS Software from the CLI

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

Step 1 Log into the switch through the console, Telnet, or SSH port of the active supervisor.

Step 2 Create a backup of your existing configuration file, if required.

Step 3 Perform the upgrade by issuing the **install all** command.

The example below demonstrates upgrading from SAN-OS 2.0(2b) to 2.1(1a) using the **install all** command with the source images located on a SCP server.

Tip Always carefully read the output of **install all**'s compatibility check. This tells you exactly what needs to be upgraded (BIOS, loader, firmware) and what modules are not hitless. If there are any questions or concerns about the results of the output, select '**n**' to stop the installation and contact the next level of support.

```
ca-9506# install all system scp://testuser@dino/tftpboot/rel/qa/2_1_1a/final/m95
00-sf1ek9-mz.2.1.1a.bin kickstart scp://testuser@dino/tftpboot/rel/qa/2_1_1a/fin
al/m9500-sf1ek9-kickstart-mz.2.1.1a.bin
For scp://testuser@dino, please enter password:
For scp://testuser@dino, please enter password:

Copying image from scp://testuser@dino/tftpboot/rel/qa/2_1_1a/final/m9500-sf1ek9
-kickstart-mz.2.1.1a.bin to bootflash:///m9500-sf1ek9-kickstart-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Copying image from scp://testuser@dino/tftpboot/rel/qa/2_1_1a/final/m9500-sf1ek9
-mz.2.1.1a.bin to bootflash:///m9500-sf1ek9-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-kickstart-mz.2.1.1a.bin
[####################] 100% -- SUCCESS
```

```
Verifying image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Extracting "svclc" version from image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:///m9500-sf1ek9-kickstart-mz
.2.1.1a.bin.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sf1ek9-kickstart-mz.2.
1.1a.bin.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes  non-disruptive       rolling
     2       yes  non-disruptive       rolling
     3       yes      disruptive       rolling  Hitless upgrade is not supported
     4       yes      disruptive       rolling  Hitless upgrade is not supported
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset



Images will be upgraded according to following table:
Module       Image       Running-Version           New-Version  Upg-Required
------  ----------  --------------------  --------------------  ------------
     1         slc                2.0(2b)               2.1(1a)           yes
     1        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     2         slc                2.0(2b)               2.1(1a)           yes
     2        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     3         ips                2.0(2b)               2.1(1a)           yes
     3        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     4       svclc                2.0(2b)               2.1(1a)           yes
     4       svcsb                1.3(5m)               1.3(5m)            no
     4       svcsb                1.3(5m)               1.3(5m)            no
     4        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     5      system                2.0(2b)               2.1(1a)           yes
     5   kickstart                2.0(2b)               2.1(1a)           yes
     5        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     5      loader                 1.2(2)                1.2(2)            no
     6      system                2.0(2b)               2.1(1a)           yes
     6   kickstart                2.0(2b)               2.1(1a)           yes
     6        bios      v1.1.0(10/24/03)      v1.1.0(10/24/03)            no
     6      loader                 1.2(2)                1.2(2)            no

Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Syncing image bootflash:///m9500-sf1ek9-kickstart-mz.2.1.1a.bin to standby.
```

```
[####################] 100% -- SUCCESS

Syncing image bootflash:///m9500-sf1ek9-mz.2.1.1a.bin to standby.
[####################] 100% -- SUCCESS

Setting boot variables.
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 5: Waiting for module online.
2005 May 20 15:46:03 ca-9506 %KERN-2-SYSTEM_MSG: mts: HA communication with standby
terminated. Please check the standby supervisor.
 -- SUCCESS

"Switching over onto standby".
```

Step 4 Exit the switch console and open a new terminal session to view the upgraded supervisor module using the **show module** command.

If the configuration meets all guidelines when the **install all** command is issued, all modules (supervisor and switching) are upgraded. This is true for any switch in the Cisco MDS 9000 Family.

Troubleshooting Cisco SAN-OS Software System Reboots

This section lists possible problems and solutions for software reboots and includes the following topics:

Power On or Switch Reboot Hangs

Symptom Power on or switch reboot hangs.

Table 2-5 Power-on or Switch Reboot Hangs

Problem	Possible Cause	Solution
Power on or switch reboot hangs for dual supervisor configuration.	The bootflash is corrupted.	See the "Recovery for Switches with Dual Supervisor Modules" section on page 2-21.
Power on or switch reboot hangs for single supervisor configuration.	The loader is corrupted.	Interrupt the boot process and reconfigure the BIOS through the console port to load a new kickstart image that updates to BIOS image. See the "Recovery Using BIOS Setup" section on page 2-15.
	The BIOS is corrupted.	Replace this module. Contact your customer support representative to return the failed module.
	The kickstart image is corrupted.	Interrupt the boot process at the >loader prompt. Update the kickstart image. See the "Recovery from the loader> Prompt" section on page 2-19.
	Boot parameters are incorrect.	Verify and correct the boot parameters and reboot.
	The system image is corrupted.	Interrupt the boot process at the switch#boot prompt. Update the system image. See the "Recovery from the switch(boot)# Prompt" section on page 2-20.

Corrupted Bootflash Recovery

All switch configurations reside in the internal bootflash. If you have a corrupted internal bootflash you could potentially lose your configuration. Be sure to save and back up your configuration files periodically. The regular switch boot goes through the following sequence (see Figure 2-2):

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

Figure 2-2 Regular Boot Sequence

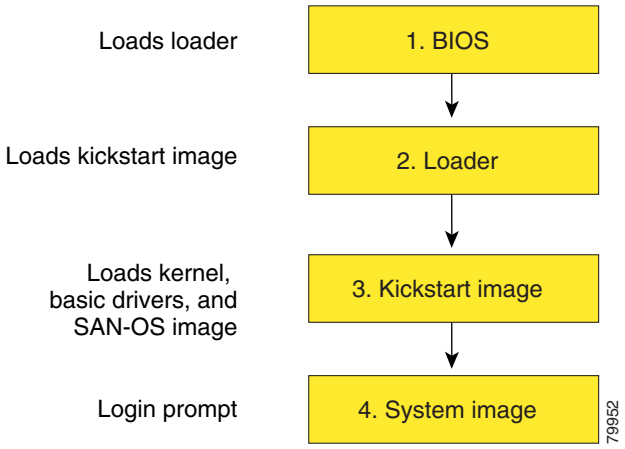

If the images on your switch are corrupted and you cannot proceed (error state), you can interrupt the switch boot sequence and recover the image by entering the BIOS configuration utility described in the following section. Access this utility only when needed to recover a corrupted internal disk.

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

Recovery procedures require the regular sequence to be interrupted. The internal switch sequence goes through four phases between the time you turn the switch on and the time the switch prompt appears on your terminal—BIOS, boot loader, kickstart, and system (see Table 2-6 and Figure 2-3).

Table 2-6 Recovery Interruption

Phase	Normal Prompt[1]	Recovery Prompt[2]	Description
BIOS	`loader>`	`No bootable device`	The BIOS begins the power-on self test, memory test, and other operating system applications. While the test is in progress, press **Ctrl-C** to enter the BIOS configuration utility and use the netboot option.
Boot loader	`Starting kickstart`	`loader>`	The boot loader uncompresses loaded software to boot an image using its file name as reference. These images are made available through bootflash. When the memory test is over, press **Esc** to enter the boot loader prompt.
Kickstart	`Uncompressing system`	`switch(boot)#`	When the boot loader phase is over, press **Ctrl-]** [3](Control key plus right bracket key) to enter the `switch(boot)#` prompt. If the corruption causes the console to stop at this prompt, copy the system image and reboot the switch.
System	`Login:`	—	Thesystem image loads the configuration file of the last saved running configuration and returns a switch login prompt.

1. This prompt or message appears at the end of each phase.
2. This prompt or message appears when the switch cannot progress to the next phase.
3. Depending on your Telnet client, these keys may be reserved and you need to remap the keystroke. Refer to the documentation provided by your Telnet client.

Figure 2-3 Regular and Recovery Sequence

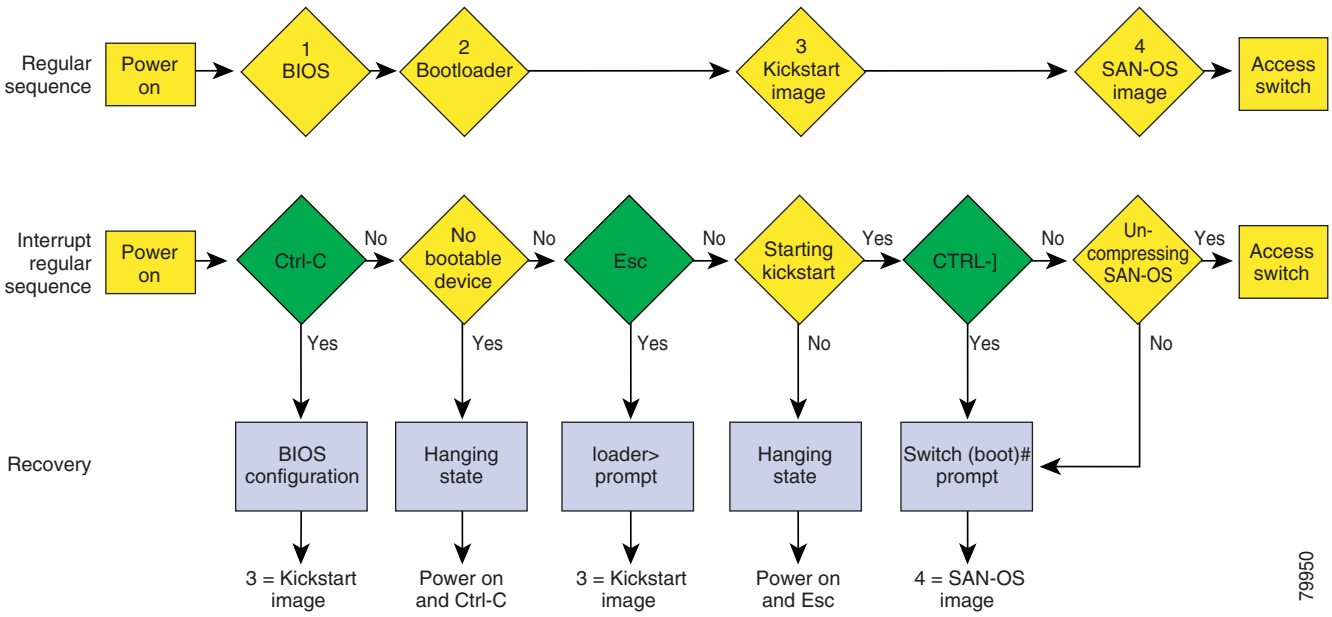

Recovery Using BIOS Setup

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

Step 1 Connect to the console port of the required switch.

Step 2 Boot or reboot the switch.

Step 3 Press **Ctrl-C** to interrupt the BIOS setup during the BIOS memory test.

You see the netboot BIOS Setup Utility screen (see Figure 2-4).

Figure 2-4 BIOS Setup Utility

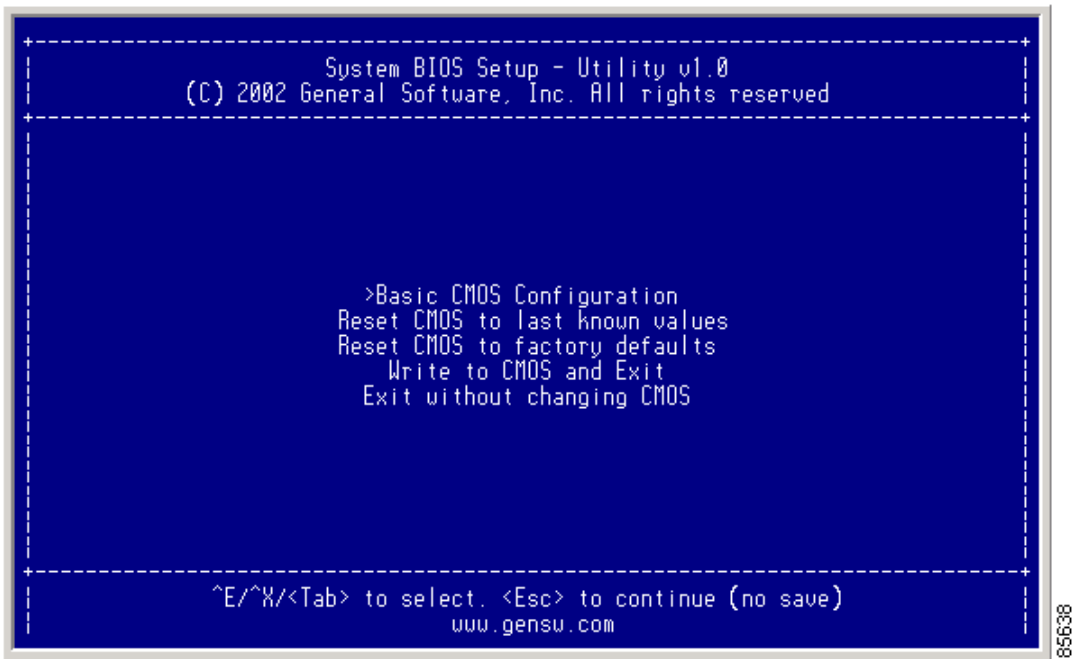

Note Your navigating options are provided at the bottom of the screen.
Tab = Jump to next field
Ctrl-E = Down arrow
Ctrl-X = Up arrow
Ctrl-H = Erase (Backspace might not work if your terminal is not configured properly.)

Step 4 Press the **Tab** key to select the Basic CMOS Configuration**.**

You see the System BIOS Setup - Basic CMOS Configuration screen (see Figure 2-5).

Figure 2-5 BIOS Setup Configuration (CMOS)

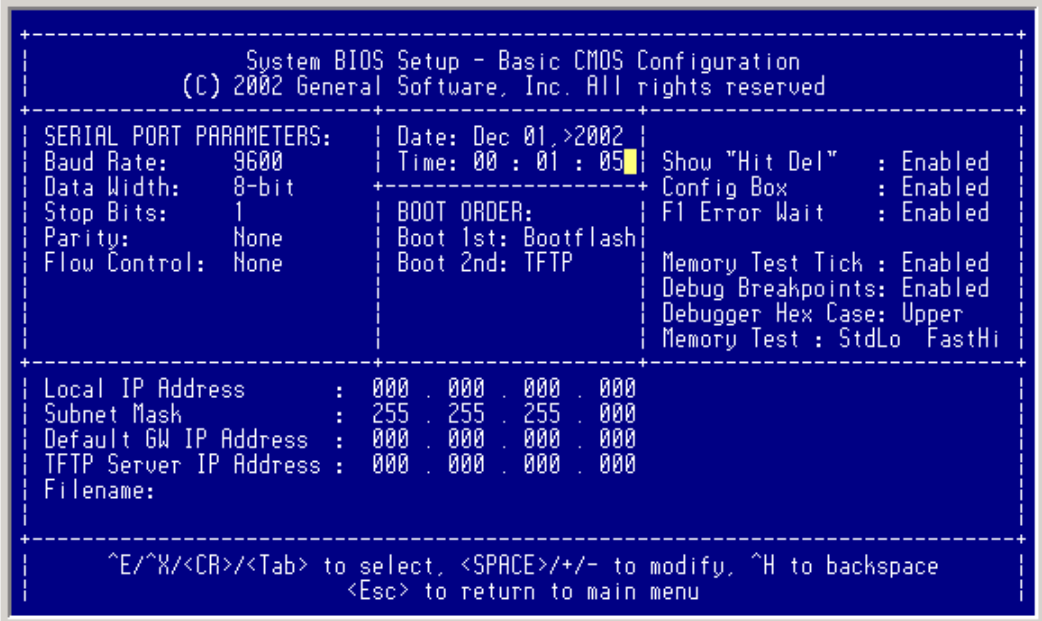

Step 5 Change the Boot 1st: field to TFTP.

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

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

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

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

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

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

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

You see the configured changes (see Figure 2-6).

Figure 2-6 BIOS Setup Configuration (CMOS) Changes

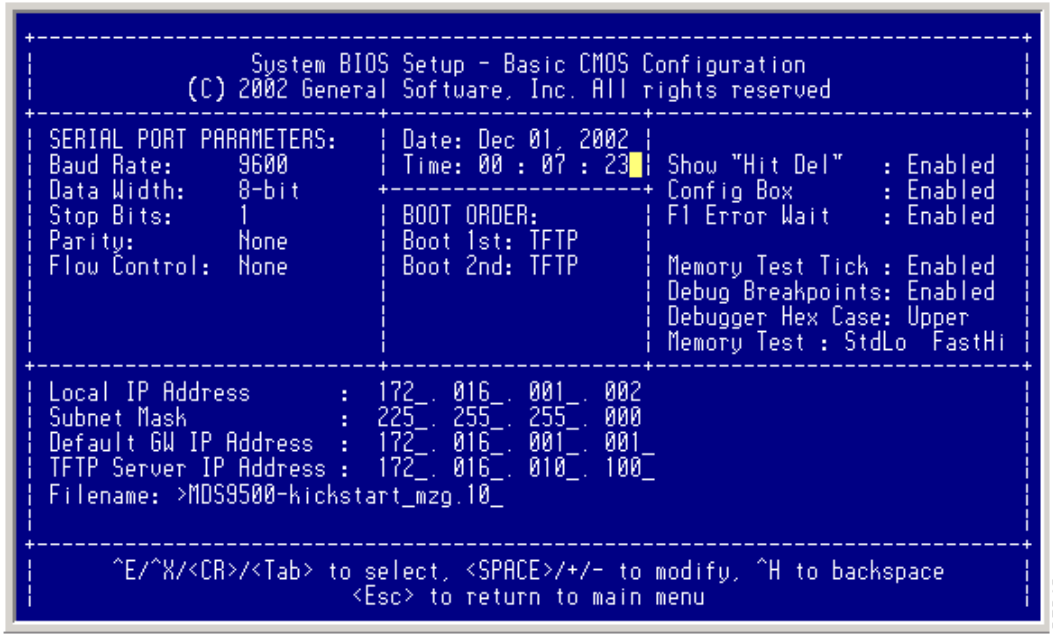

Step 12 Press the **Esc** key to return to the main menu.

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

Note These changes are saved in the CMOS.

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

You see the following prompt:

```
switch(boot)#
```

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

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

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

Step 15 Follow the procedure specified in the "Recovery from the switch(boot)# Prompt" section on page 2-20.

Recovery from the loader> Prompt

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

Tip Use the **help** command at the `loader>` prompt to display a list of commands available at this prompt or to obtain more information about a specific command in that list.

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

Step 1 Enter the local IP address and the subnet mask for the switch at he `loader>` prompt, and press **Enter**.

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

Step 2 Specify the IP address of the default gateway.

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

Step 3 Boot the kickstart image file from the required server.

```
loader> boot tftp://172.16.10.100/kickstart-image1
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Booting: /kick-282 console=ttyS0,9600n8nn quiet loader_ver= "2.1(2)"....
.........................................Image verification OK
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
switch(boot)#
```

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

Step 4 Issue the **init system** command at the `switch(boot)#` prompt.

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

Step 5 Follow the procedure specified in the "Recovery from the switch(boot)# Prompt" section on page 2-20.

Recovery from the switch(boot)# Prompt

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

Step 1 Change to configuration mode and configure the IP address of the mgmt0 interface.

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

Step 2 Follow this step if you issued an **init system** command. Otherwise, skip to Step 3.

a. Issue the **ip address** command to configure the local IP address and the subnet mask for the switch.

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

b. Issue the **ip default-gateway** command to configure the IP address of the default gateway.

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

Step 3 Issue the **no shutdown** command to enable the mgmt0 interface on the switch.

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

Step 4 Enter **end** to exit to EXEC mode.

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

Step 5 If you believe there are file system problems, issue the init **system check-filesystem** command. As of Cisco MDS SAN-OS Release 2.1(1a), this command checks all the internal file systems and fixes any errors that are encountered. This command takes considerable time to complete.

```
switch(boot)# init system check-filesytem
```

Step 6 Copy the system image from the required TFTP server.

```
switch(boot)# copy tftp://172.16.10.100/system-image1 bootflash:system-image1
```

Step 7 Copy the kickstart image from the required TFTP server.

```
switch(boot)# copy tftp://172.16.10.100/kickstart-image1 bootflash:kickstart-image1
```

Step 8 Verify that the system and kickstart image files are copied to your bootflash: file system.

```
switch(boot)# dir bootflash:
12456448     Jul 30 23:05:28 1980  kickstart-image1
12288        Jun 23 14:58:44 1980  lost+found/
27602159     Jul 30 23:05:16 1980  system-image1

Usage for bootflash://sup-local
  135404544 bytes used
   49155072 bytes free
  184559616 bytes total
```

Step 9 Load the system image from the bootflash: files system.

```
switch(boot)# load bootflash:system-image1
Uncompressing system image: bootflash:/system-image1
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC

Would you like to enter the initial configuration mode? (yes/no): yes
```

Note If you enter **no** at this point, you will return to the `switch#` login prompt, and you must manually configure the switch.

Recovery for Switches with Dual Supervisor Modules

This section describes how to recover when one or both supervisor modules in a dual supervisor switch have corrupted bootflash.

Recovering One Supervisor Module With Corrupted Bootflash

If one supervisor module has functioning bootflash and the other has corrupted bootflash, follow these steps:

Step 1 Boot the functioning supervisor module and log on to the switch.

Step 2 At the `switch#` prompt on the booted supervisor module, issue the **reload module** *slot* **force-dnld** command, where *slot* is the slot number of the supervisor module with the corrupted bootflash.

The supervisor module with the corrupted bootflash performs a netboot and checks the bootflash for corruption. When the bootup scripts discovers that the bootflash is corrupted, it performs an **init system**, which fixes the corrupt bootflash. The supervisor boots up as the HA Standby.

Recovering Both Supervisor Modules With Corrupted Bootflash

If both supervisor modules have corrupted bootflash, follow these steps:

Step 1 Boot up the switch and press the **Esc** key after the BIOS memory test to interrupt the boot loader.

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

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

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

You see the `loader>` prompt.

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

Tip Use the **help** command at the `loader>` prompt to display a list of commands available at this prompt or to obtain more information about a specific command in that list.

Step 2 Specify the local IP address and the subnet mask for the switch.

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

Step 3 Specify the IP address of the default gateway.

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

Step 4 Boot the kickstart image file from the required server.

```
loader> boot tftp://172.16.10.100/kickstart-latest
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Booting: /kick-282 console=ttyS0,9600n8nn quiet loader_ver= "2.1(2)"....
.........................................Image verification OK
Starting kernel...
```

```
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
switch(boot)#
```

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

Step 5 Issue the **init-system** command to repartition and format the bootflash.

Step 6 Perform the procedure specified in the "Recovery from the switch(boot)# Prompt" section on page 2-20.

Step 7 Perform the procedure specified in the "Recovering One Supervisor Module With Corrupted Bootflash" section on page 2-21 to recover the other supervisor module.

Note If you do not issue the **reload module** command when a boot failure has occurred, the active supervisor module automatically reloads the standby supervisor module within 3 to 6 minutes after the failure.

Recognizing Error States

If you see the error messages displayed in Figure 2-7 or Figure 2-8, follow the procedure specified in the "Recovery Using BIOS Setup" section on page 2-15.

Figure 2-7 Error State if Powered On and Ctrl-C Is Entered

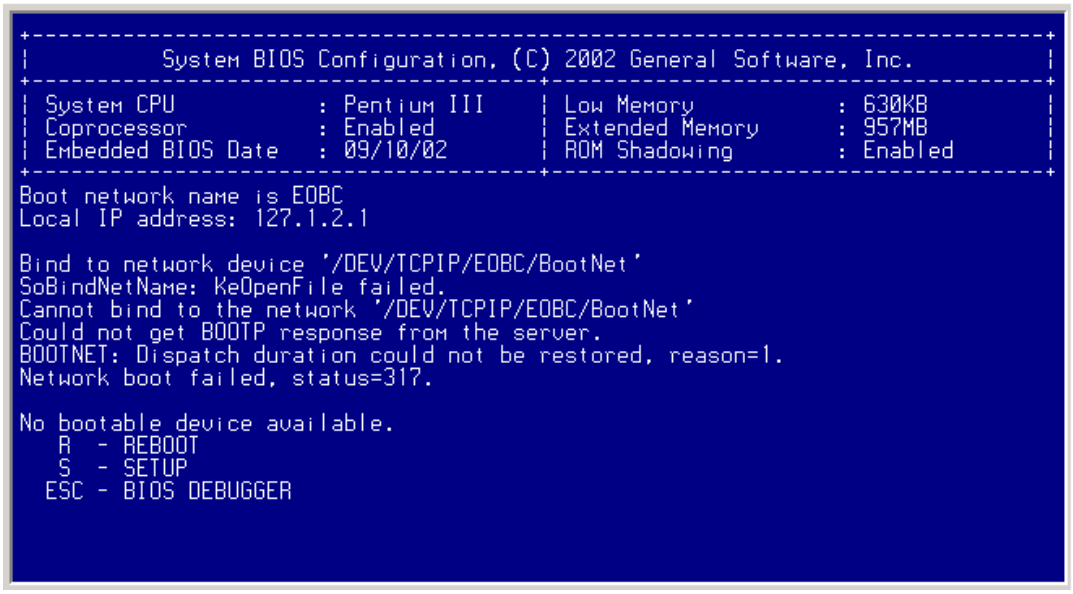

Figure 2-8 Error State if Powered On and Esc Is Pressed

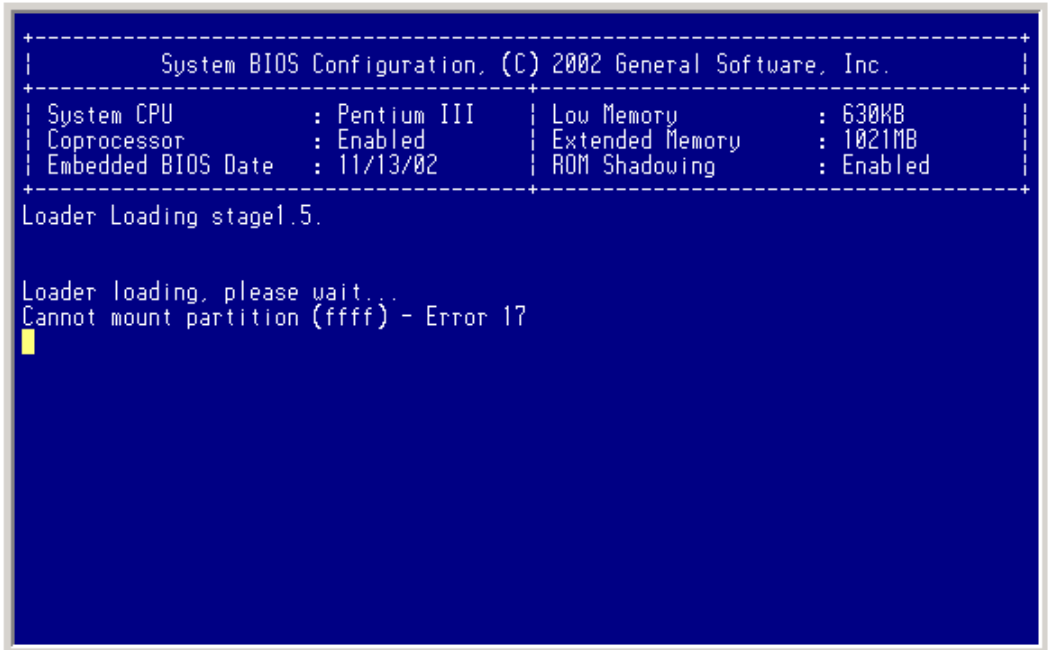

Switch or Process Resets

When a recoverable or nonrecoverable error occurs, the switch or a process on the switch may reset.

Symptom The switch or a process on the switch reset.

Table 2-7 Switch or Process Resets

Problem	Possible Cause	Solution
The switch or a process on the switch resets.	A recoverable error occurred on the system or on a process in the system.	Cisco SAN-OS automatically recovered from the problem. See the "Recoverable System Restarts" section on page 2-25 and the "Switch or Process Resets" section on page 2-24.
	A nonrecoverable error occurred on the system.	Cisco SAN-OS cannot recover automatically from the problem. See the "Unrecoverable System Restarts" section on page 2-29 to determine the cause.
	A clock module failed.	Verify that a clock module failed. See the "Troubleshooting Clock Module Issues" section on page 3-13. Replace the failed clock module during the next maintenance window.

Recoverable System 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# show log logfile | include error
```

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

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 I/O)
- R = runnable (on run queue)
- S = sleeping
- T = traced or stopped
- Z = defunct (“zombie”) process
- NR = notrunning
- ER = should be running but currently notrunning

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 example. (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     -  bdflush
    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
snsm              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
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 7FFFFC9C         EBP 7FFFFCAC
    EAX 00000008         XDS 0000002B         XES 0000002B
    EAX 00000003 (orig)  EIP 2ABF5EF4         XCS 00000023
    EFL 00000246         ESP 7FFFFC5C         XSS 0000002B
Stack: 128 bytes. ESP 7FFFFC5C, TOP 7FFFFD10
0x7FFFFC5C: 0804F990 0804C416 00000003 0804E994 ................
0x7FFFFC6C: 00000008 0804BF95 2AC451E0 2AAC24A4 .........Q.*.$.*
0x7FFFFC7C: 7FFFFD14 2AC2C581 0804E6BC 7FFFFCA8 .......*........
```

```
0x7FFFFC8C: 7FFFFC94 00000003 00000001 00000003 ................
0x7FFFFC9C: 00000001 00000000 00000068 00000000 ........h.......
0x7FFFFCAC: 7FFFFCE8 2AB4F819 00000001 7FFFFD14 .......*........
0x7FFFFCBC: 7FFFFD1C 0804C470 00000000 7FFFFCE8 ....p...........
0x7FFFFCCC: 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# show system 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**, c**ompare 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
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 module-num column shows the slot number on which the core was generated. In the previous example, 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 module 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# show 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
    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: 2AC8DFF8 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 to configure the switch to use TFTP to send the core dump to a TFTP server.

system 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 interval is not configurable.
- The copy of a specific core file to a TFTP server 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 Determine the cause and resolution for the restart condition by contacting your customer support representative and asking them to review your core dump.

See also the "Troubleshooting Supervisor Issues" section on page 3-15 or the "Troubleshooting Switching and Services Modules" section on page 3-22.

Unrecoverable System Restarts

An unrecoverable system restart might occur in the following cases:

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

To respond to an unrecoverable reset, see the "Troubleshooting Cisco SAN-OS Software System Reboots" section on page 2-12.

The **show system reset-reason** CLI command displays the following information:

- In a Cisco MDS 9500 Series switch, the last four reset-reason codes for the supervisor module in slot 5 and slot 6 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9200 Series switch, the last four reset-reason codes for the supervisor module in slot 1 are displayed.
- The **show system reset-reason module number** command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module are not displayed.
- Find the overall history of when and why expected and unexpected reloads occur.
- Timestamp of when the reset or reload occurred
- Reason for the reset or reload of a module
- The service that caused the reset or reload (not always available)
- The software version that was running at the time of the reset or reload

Example 2-2 show system reason-reset Command Output

```
switch# show system reset-reason module 5
----- reset reason for module 5 -----
1) At 224801 usecs after Fri Jan 21 16:36:40 2005
Reason: Reset Requested by CLI command reload
Service:
Version: 2.1(2)
2) At 922828 usecs after Fri Jan 21 16:02:48 2005
Reason: Reset Requested by CLI command reload
Service:
Version: 2.1(2)
3) At 318034 usecs after Fri Jan 21 14:03:36 2005
Reason: Reset Requested by CLI command reload
Service:
Version:2.1(2)
4) At 255842 usecs after Wed Jan 19 00:07:49 2005
Reason: Reset Requested by CLI command reload
Service:
Version: 2.1(2)
```

Recovering the Administrator Password

You can access the switch if you forget the administrator password by following the directions in Table 2-8.

Symptom You forgot the administrator password for accessing a switch.

Table 2-8 Recovering Administrator Password

Problem	Solution
You forgot 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, refer to the Cisco MDS 9000 Family Configuration Guide at the following website: http://cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html

Miscellaneous Software Image Issues

This section includes software image issues reported by the relevant release notes and includes the following topics:

All Ports Down Because of System Health Failure

Symptom Console reports all ports on a module are down because of a system health failure.

Table 2-9 All Ports are Down Because of a System Health Failure.

Symptom	Possible Cause	Solution
The system console reports that the module's ports are down because of to a system health failure.	An incorrect device instance on the Cisco MDS 9000 modules might get reinitialized from an error recovery mechanism, leaving the module in an unusable state. In some cases, the module may reboot.	Downgrade to a Cisco SAN-OS Release 2.0(x) version supported by your OSM.
		Upgrade to Cisco SAN-OS Release 2.1.2 or 2.1(1b). Resetting the module will clear the problem, but the problem could reoccur unless you are using a SAN-OS version with the bug fix.

Switch Reboots after FCIP Reload

Symptom Switch rebooted after FCIP module was reloaded, upgraded or downgraded.

Table 2-10 Switch Reboot after FCIP Reload

Symptom	Possible Cause	Solution
Switch rebooted after FCIP module was reloaded, upgraded, or downgraded	If an IPS module with operational FCIP PortChannels is reloaded, upgraded, or downgraded, the supervisor module may be reloaded causing the system to reboot.	Before reloading, upgrading, or downgrading an IPS module, shut down all FCIP PortChannels on the module.

FCIP Link Fails to Come Up

Symptom A newly configured FCIP link may fail to come up when running on an MPS-14/2 module.

Table 2-11 FCIP Link Fails to Come Up

Symptom	Possible Cause	Solution
A newly configured FCIP link may fail to come up when running on an MPS-14/2 module.	This symptom may occur following an upgrade from Cisco MDS SAN-OS Release 2.0(1b) to Release 2.0(3) and the configuration of a new FCIP link.	Reload the MPS-14/2 module using the **reload module** *module-number* command, where *module-number* is a specific module.

Cannot Create, Modify, or Delete Admin Role

Symptom Cannot create, modify, or delete the admin role.

Table 2-12 Cannot Create, Modify, or Delete Admin Role

Symptom	Possible Cause	Solution
Cannot create, modify, or delete the admin role	After upgrading to Cisco SAN-OS Release 2.0, it is no longer possible to create, modify, or delete the admin role.	Create the admin role before upgrading to Cisco SAN-OS Release 2.0.

FC IDs Change after Link Reset

Symptom FC IDs change after a link resets.

Table 2-13 FC IDs Change After a Link Reset

Symptom	Possible Cause	Solution
FC IDs change after a link resets.	Following an upgrade from Cisco SAN-OS Release 1.1 to Cisco SAN-OS Release 1.3 or later, with persistent FC ID enabled, the FC IDs for the storage arrays may get changed after a link flap.	Reconfigure the FC IDs as necessary.

Switch Displays Wrong User

Symptom Switch displays the wrong user with the s**how running**-config CLI command.

Table 2-14 Switch Displays Wrong User

Symptom	Possible Cause	Solution
Switch displays the wrong user with the s**how running**-config CLI command.	When you perform a nondisruptive upgrade from Cisco SAN-OS Release 1.3(x) to Cisco SAN-OS Release 2.0(x), and then issue the **show running-config** command, the switch displays the wrong user. The user shown will be inconsistent with the user shown when you issue the **show user-account** command.	Recreate the user.

CHAPTER 3

Troubleshooting Hardware

This chapter describes how to identify and resolve problems that might occur in the hardware components of the Cisco MDS 9000 Family. It includes the following sections:

Overview

The key to success when troubleshooting the system hardware is to isolate the problem to a specific system component. The first step is to compare what the system is doing to what it should be doing. Because a startup problem can usually be attributed to a single component, it is more efficient to isolate the problem to a subsystem rather than troubleshoot each separate component in the system.

Problems with the initial power up are often caused by a module that is not firmly connected to the backplane or a power supply that has been disconnected from the power cord connector.

Overheating can also cause problems with the system, though typically only after the system has been operating for an extended period of time. The most common cause of overheating is the failure of a fan module.

The Cisco MDS 9000 Family includes the following subsystems on most chassis:

- Power supply— This includes the power supply fans.
- Fan module—The chassis fan module should operate whenever system power is on. You should see the Fan LED turn green and should hear the fan module to determine whether or not it is operating. If the Fan LED is red, this indicates that one or more fans in the fan module is not operating. You

should immediately contact your customer service representative (see the "Steps to Perform Before Calling TAC" section on page A-1). There are no installation adjustments that you can make if the fan module does not function properly at initial startup.

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

- Supervisor module—The supervisor module contains the operating system software, so check your supervisor module if you have trouble with the system software. Status LEDs on the supervisor module indicate whether or not the supervisor module can initialize a switching or services module.

  If you have a redundant supervisor module, refer to the following website for the latest Cisco MDS 9000 Family configuration guides for descriptions of how the redundant supervisor module comes online and how the software images are handled: http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm.
- Switching or services module—Status LEDs on each module indicate if it has been initialized by the supervisor module. A module that is partially installed in the backplane can cause the system to halt.

Best Practices

You should consider the best practices recommended in this section to ensure the proper installation, initialization, and operation of your switch. This section includes the following topics:

Best Practices for Switch Installation

Follow these best practices when installing your switch:

- Plan your site configuration and prepare the site before installing the chassis.
- Verify that you have the appropriate power supplies for your chassis configuration.
- Install the chassis following the rack and airflow guidelines presented in the associated Cisco MDS 9000 Family hardware installation guide for your chassis.
- Verify that the chassis is adequately grounded.

Best Practices for System Initialization

When the initial system boot is complete, verify the following:

- Power supplies are supplying power to the system. See the "Troubleshooting Power Supply Issues" section on page 3-4.
- The system fan module is operating. See the "Troubleshooting Fan Issues" section on page 3-9.

- The system software boots successfully. Refer to the following website for the latest Cisco MDS 9000 Family configuration guides containing information on booting the system and initial configuration tasks: http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm.
- The supervisor module and all switching or services modules are installed correctly and each one initialized without problems. See the "Troubleshooting Supervisor Issues" section on page 3-15.

If all of these conditions are met and the hardware installation is complete, see the rest of this document to troubleshoot any other software issues.

If any of these conditions are not met, use the procedures in this chapter to isolate and, if possible, resolve the problem.

Best Practices for Supervisor Modules

As a best practice, we recommend that you take the following actions to ensure proper operation of your supervisor modules:

- Make sure both supervisors have their Flash memory loaded with the same versions of kickstart and system images.
- Make sure that the proper boot statements for the active and standby supervisors are set to run the same code.
- Once the boot statements are configured on the active supervisor, issue the **copy running-config startup-config** command.
- Make a copy of the running configuration to CompactFlash for a safe backup.
- Always issue the **copy running-config startup-config** CLI command when modifying the running configuration and you have ensured that the system is operating properly.
- Never use the **init system** CLI command unless you understand that you will lose the running and startup configuration as well as all files stored on bootflash:.
- Keep backup copies of running kickstart and system images on CompactFlash.

Troubleshooting Startup Issues

LEDs indicate all system states in the startup sequence. By checking the LEDs, you can determine when and where the system failed in the startup sequence.

To identify startup problems, follow these steps:

Step 1 Turn on the power supplies by turning or pressing the switch on (|). You should immediately hear the system fan module begin to operate. If not, see the "Troubleshooting Power Supply Issues" section on page 3-4.

Step 2 If you determine that the power supplies are functioning normally and the fan module is faulty, see the "Troubleshooting Fan Issues" section on page 3-9.

Step 3 Verify that the LEDs on the supervisor module display as follows:

a. The Status LED flashes orange once and stays orange during diagnostic boot tests. It turns green when the module is operational (online). If the system software cannot start up, this LED stays orange.

b. The System LED turns green, indicating that all chassis environmental monitors are reporting that the system is operational. If one or more environmental monitors reports a problem, the System LED is orange or red.

c. The Active LED turns green, indicating that the supervisor module is operational and active. If the supervisor module is in standby mode, the Active LED is orange.

d. Each Link LED flashes orange once and stays orange during diagnostic boot tests, and turns green when the module is operational (online). If no signal is detected, the Link LED turns off. The link LED blinks orange if the port is bad.

If any LEDs on the supervisor module front panel are red or orange after the initialization time, see the "Troubleshooting Supervisor Issues" section on page 3-15. If you have a redundant supervisor module, refer to the following website for the latest Cisco MDS 9000 Family configuration guides for descriptions of the supervisor module LEDS, how the redundant supervisor module comes online, and how the software images are handled:

http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm.

Step 4 Verify that the Status LEDs on the supervisor module and on each switching or services module are green when the supervisor module completes initialization. This LED indicates that the modules are receiving power, have been recognized by the supervisor module, and contain a valid Flash code version. This LED does not indicate the state of the individual interfaces on the switching modules. If a Status LED is red or orange, see the "Troubleshooting Supervisor Issues" section on page 3-15.

Step 5 Verify that the terminal is set correctly and that it is connected properly to the supervisor module console port if the boot information and system banner are not displayed.

Troubleshooting Power Supply Issues

This section describes power supply problems and includes the following topics:

All Power Supply LEDS Are Off

Symptom All power supply LEDS are off.

The following system messages may be generated with this symptom:

Error Message `PLATFORM-2-PS_FAIL: Power supply [dec] failed or shutdown (Serial No. [chars]).`

Explanation Power supply failed or has been shut down.

Recommended Action Enter the **show environment power** and **show platform internal info** CLI commands or similar Fabric Manager or Device Manager command to collect more information. Refer to power supply documentation in the relevant hardware installation guide to learn more on increasing or decreasing power supply capacity and configuring power supplies.

Error Message `PLATFORM-2-PS_MISMATCH: Detected power supply [chars]. This reduces the redundant power available to the system and can cause service disruptions (Serial No. [chars]).`

Explanation Detected a new power supply that has reduced capacity compared to an existing power supply.

Recommended Action Refer to power supply document on increasing decreasing power supply capacity and configuring power supplies. Enter the **show environment power** and **show platform internal info** CLI command or similar Fabric Manager/Device Manager command to collect more information.

Error Message `PLATFORM-5-PS_REMOVE: Power supply [dec] removed (Serial No. [chars]).`

Explanation Power supply has been removed.

Recommended Action No action is required.

Table 3-1 All Power Supply LEDS Are Off

Symptom	Possible Cause	Solution
All power supply LEDS are off.	Power supply is not correctly seated in the chassis.	Remove and reinstall the power supply. Refer to the appropriate hardware installation guide for your chassis.
	Power supply is shut down.	Choose **Physical > Power Supplies** and check the OperStatus on Device Manager, or use the **show environment power** CLI command to determine if the power supply is shut down. If the status is shutdown, then the supervisor has shutdown the power supply. The supervisor shuts down the lower capacity power supply only if it detects a mismatched pair of power supplies and the mode is redundant or there is a transition from combined to redundant mode. If both power supplies are the same capacity or the mode is combined, Cisco SAN-OS never shuts down a power supply.
	Power supply is not operational.	Troubleshoot the power supplies. See the "Troubleshooting the Power Supplies" section on page 3-8.

Power Supply Input Ok LED is Red

Symptom Power supply Input Ok LED is red.

Table 3-2 Power Supply INput Ok LED Is Red

Symptom	Possible Cause	Solution
Power supply Input Ok LED is red.	Power supply is not correctly seated in the chassis.	Remove and reinstall the power supply. Refer to the appropriate hardware installation guide for your chassis.
	PEMs on a Cisco MDS 9500Series chassis are not correctly installed.	Remove and reinstall the power supply PEMs. Refer to the appropriate hardware installation guide for your chassis.
	External power source is not operational.	Power down the switch and verify external power source. Use independent power sources to each redundant power supply in a Cisco MDS 9500 Series director.
	Power supply is not operational.	Troubleshoot the power supplies. See the "Troubleshooting the Power Supplies" section on page 3-8.

Power Supply Output Failed LED is On

Symptom Power Supply Output Failed LED is on.

Table 3-3 Power Supply Output Failed LED is On

Symptom	Possible Causes	Solutions
Power Supply Output Failed LED is on.	Power supply is not operational.	Troubleshoot the power supplies. See the "Troubleshooting the Power Supplies" section on page 3-8.

Power Supply Fan Ok LED is Red

Symptom Power supply Fan Ok LED is red.

The following system messages may be generated with this symptom:

Error Message `PLATFORM-2-PS_FANFAIL: Fan in Power supply [dec] failed.`

Explanation Fan module in the power supply has failed.

Recommended Action Enter the **show environment power** and **show platform internal info** CLI command or similar Fabric Manager/Device Manager command to collect more information.

Introduced Cisco MDS SAN-OS Release 1.3(1).

Table 3-4 Power Supply Fan Ok LED is Red

Symptom	Possible Cause	Solution
Power supply Fan Ok LED is red.	Fan has failed on the power supply.	Choose **Physical > Temperature** sensors on Device Manager or use the **show environment temperature** CLI command to verify that the chassis temperature is normal. Verify that no temperature sensors are approaching the minor thresholds. If the temperature sensors are near or over a threshold value, you should replace the power supply.
	Power supply is not operational.	Troubleshoot the power supplies. See the "Troubleshooting the Power Supplies" section on page 3-8.

Troubleshooting the Power Supplies

To isolate a power supply problem, follow these steps:

Step 1 Verify that the Input Ok LED on the power supply is green. If the Input Ok LED is green, the AC or DC source is good and the power supply is functional.

Step 2 If the Input Ok LED is off, first ensure that the power supply is flush with the chassis. Turn the power switch off, tighten the captive screw(s), and then turn the power switch on (|). If the Input Ok LED remains off, there might be a problem with the AC source or the DC source, or the power cable.

- **a.** Turn off the power to the switch by pressing or turning both power switches to 0, connect the power cord to another power source if one is available, and turn the power on. If the Input Ok LED is now green, the problem was the first power source.
- **b.** If the Input Ok LED fails to light after you connect the power supply to a new power source, replace the power cord and turn the switch on. If the Input Ok LED lights at this point, return the first power cord for replacement.
- **c.** If the Input Ok LED still fails to light when the switch is connected to a different power source with a new power cord, the power supply is probably faulty. If a second power supply is available, install it in the second power supply bay and contact your customer service representative for further instructions.

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

Step 3 Repeat Step 1 if you have a second (redundant) power supply.

Step 4 Choose **Physical > Power Supplies** on Device Manager or use the **show environment power** CLI command to verify the status of your power supplies. (See Example 3-1.)

Example 3-1 Output of show environment power

```
switch# show environment power
-------------------------------------------------------
PS  Model                Power     Power       Status
                         (Watts)   (Amp @42V)
-------------------------------------------------------
1   DS-CAC-1900W         1019.34   24.27        ok
2   DS-CAC-1900W         1019.34   24.27        ok


Mod Model                  Power     Power      Power     Power       Status
                         Requested Requested   Allocated Allocated
                          (Watts)   (Amp @42V)  (Watts)   (Amp @42V)
--- -------------------  --------  ----------  --------- ----------  ----------
3    DS-X9016            220.08      5.24       220.08     5.24      powered-up
4    DS-X9308-SMIP       210.00      5.00       210.00     5.00      powered-up
5    DS-X9530-SF1-K9     220.08      5.24       220.08     5.24      powered-up


Power Usage Summary:
------------------
Power Supply redundancy mode:                redundant

Total Power Capacity                         1019.34  W
```

```
Power reserved for Supervisor(s)[-]           440.16   W
Power reserved for Fan Module(s)[-]           126.00   W
Power currently used by Modules[-]            430.08   W
```

If you are unable to resolve the problem or if you determine that either a power supply or backplane connector is faulty, contact your customer support representative.

Troubleshooting Fan Issues

This section describes fan failure problems and includes the following topics:

Fan Is Not Spinning

Symptom Fan is not spinning.

Table 3-5 Fan Is Not Spinning

Symptom	Possible Cause	Solution
Fan is not spinning.	Fan is not correctly seated in the chassis.	Loosen the captive screws, remove the fan module and reinstall it to ensure that the fan module is seated properly. Tighten all captive screws, and then restart the system.
	Power supply is not operational.	Troubleshoot the power supplies. See the "Troubleshooting Power Supply Issues" section on page 3-4.

Fan Is Spinning, But Fan LED is Red

Symptom Fan is spinning, but fan LED is red.

Table 3-6 Fan Is Spinning, Fan LED is Red

Symptom	Possible Cause	Solution
Fan is spinning but fan LED is red.	Fan is not correctly seated in the chassis.	Loosen the captive screws, remove the fan module and reinstall it to ensure that the fan module is seated properly. Tighten all captive screws, and then restart the system.
	Fan module has failed.	Troubleshoot the Fan Module. See the "Troubleshooting a Fan Failure Using the CLI" section on page 3-11.

Troubleshooting a Fan Failure Using Device Manager

To troubleshoot a fan module problem using Device Manager, follow these steps:

Step 1 Choose **Physical > Fan**. You see the Fan Status dialog box.

Step 2 If the OperStatus is failure, one or more fans are not operational. Replace the failed fan module before your switch overheats. You should see the following system message in the switch log:

Error Message `PLATFORM-1-CASA_FAN_FAIL: Fan module [dec] Failed.`

Explanation Fan module failed and needs to be replaced. This can lead to overheating and temperature alarms.

Recommended Action Enter the **show platform internal info** CLI command or similar Fabric Manager/Device Manager command to collect more information.

Step 3 If the OperStatus is absent, the fan module has been removed. As soon as the fan module is removed, Cisco SAN-OS starts a 5 minute countdown.

Caution If the fan module is not reinserted within 5 minutes, the entire switch is shutdown.

Software reads a byte on the SEEPROM to determine if the fan module is present. If the fan module is partially inserted or software is unable to access the SEEPROM on the fan module for any other reason, then Cisco SAN-OS cannot distinguish this case from a real fan module removal. The switch will be shut down in five minutes. The following priority 0 syslog messages are printed every five seconds:

Error Message `PLATFORM-0-FAIL_REMOVED: Fan module removed. Fan module has been absent for [dec] seconds.`

Explanation Fan module was removed. This could lead to temperature alarms.

Recommended Action Replace the fan module immediately.

Step 4 Remove and reinstall or replace the fan module. If the Fan LED is still red, the system detects a fan module failure. Contact your customer service representative for instructions.

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

Troubleshooting a Fan Failure Using the CLI

To troubleshoot a fan module problem using the CLI, follow these steps:

Step 1 Use the **show environment fan** CLI command and verify the status of each fan type. (See Example 3-2.)

Example 3-2 show environment fan Output

```
switch# show environment fan
----------------------------------------------------
Fan              Model           Hw         Status
----------------------------------------------------
Chassis         DS-9SLOT-FAN    1.2          ok
PS-1             --              --          ok
PS-2             --              --          absent
```

Step 2 If the fan status is failure, one or more fans are not operational. Replace the failed fan module before your switch overheats. You should see the following system message in the log:

Error Message `PLATFORM-1-CASA_FAN_FAIL: Fan module [dec] Failed.`

Explanation Fan module failed and needs to be replaced. This can lead to overheating and temperature alarms.

Recommended Action Enter the **show platform internal info** CLI command to collect more information.

Step 3 If the fan status is absent, the fan module has been removed. As soon as the fan module is removed, Cisco SAN-OS starts a 5 minute countdown.

Caution If the fan module is not reinserted within five minutes, the entire switch is shut down.

Software reads a byte on the SEEPROM to determine if the fan module is present. If the fan module is partially inserted or software is unable to access the SEEPROM on the fan module for any other reason, then Cisco SAN-OS cannot distinguish this case from a real fan module removal. The switch will be shut down in five minutes. The following priority 0 syslog messages are printed every five seconds:

Error Message `PLATFORM-0-FAIL_REMOVED: Fan module removed. Fan module has been absent for [dec] seconds.`

Explanation Fan module was removed. This could lead to temperature alarms.

Recommended Action Replace the fan module immediately.

Step 4 Remove and reinstall or replace the fan module. If the Fan LED is still red, the system detects a fan module failure. Contact your customer service representative for instructions.

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

Temperature Threshold Violations

Each card in the chassis has at least two temperature sensors. Each temperature sensor is configured with a minor and a major threshold. Example 3-3 gives the **show environment temperature** CLI command sample output. It shows how temperature information can be retrieved from the switch. Choose **Physical > Temperature Sensors** on Device Manager to view a similar output.

Example 3-3 Output of show environment temperature Command

```
switch# show environment temperature
-----------------------------------------------------------------
Module   Sensor   MajorThresh   MinorThres   CurTemp     Status
                  (Celsius)     (Celsius)    (Celsius)
-----------------------------------------------------------------
4        Outlet     75            60           36          ok
4        Intake     65            50           29          ok

5        Outlet     75            60           35          ok
5        Intake     65            50           34          ok

6        Outlet     75            60           35          ok
6        Intake     65            50           34          ok

9        Outlet     75            60           45          ok
9        Intake     65            50           40          ok
```

The intake sensor is placed at the airflow intake and is the most critical indicator of module temperature. All Cisco SAN-OS actions are taken when the major threshold of an intake sensor is exceeded.

A minor threshold violation or a major threshold violation on an outlet sensor results in the following system message:

Error Message `PLATFORM-0-MOD_TEMPMAJALRM: Module [dec] reported major temperature alarm.`

Explanation Module in the slot exceeded a major temperature threshold.

Recommended Action Enter the **show environment temperature** CLI command or choose **Physical > Temperature Sensors** on Device Manager to collect more information.

This also generates a Call Home event and an SNMP notification.

A major temperature threshold violation on a module intake sensor results in the following system message:

Error Message `PLATFORM-0-MOD_TEMPSHUTDOWN: Module [dec] powered down due to major temperature alarm.`

Explanation Module shutdown due to temperature exceeding major threshold.

Recommended Action Enter **show environment temperature** CLI command or similar Fabric Manager/Device Manager command to collect more information.

If Cisco SAN-OS detects a major temperature threshold violation on a redundant supervisor intake sensor, it immediately shuts down the redundant supervisor. This will result in either a switchover or the standby shutting down, depending on the supervisor that violated the threshold.

If Cisco SAN-OS detects a major temperature threshold violation on an intake sensor on the only operational supervisor in a switch, a 120 second countdown starts. If the temperature recovers, the countdown is discarded. Otherwise, the switch power supplies are shutdown. The following syslog messages are printed every five seconds during the countdown

Error Message `PLATFORM-0-SYS_RESET: [chars] System shutdown in [dec] seconds.`

Explanation System shutdown in the number of seconds shown in the error message.

Recommended Action Enter **show environment temperature** CLI command or similar Fabric Manager/Device Manager command to collect more information.

Sometimes, a temperature sensors fails. No explicit action is taken for this condition except generating the following system message:

Error Message `PLATFORM-5-MOD_TEMPFAIL: Module [dec] temperature sensor failed.`

Explanation Module contains a faulty temperature sensor.

Recommended Action Enter the **show environment temperature** CLI command or similar Fabric Manager/Device Manager command to collect more information.

Troubleshooting Clock Module Issues

A Cisco MDS 9500 Series director has two clock modules, A and B. Use the **show environment clock** CLI command to view the clock module status. (See Example 3-4.)

Example 3-4 Output of show environment clock Command

```
switch# show environment clock
-------------------------------------------------------------
Clock           Model                Hw         Status
-------------------------------------------------------------
A               DS-C9500-CL          0.0        ok/active
B               DS-C9500-CL          0.0        ok/standby
```

On a clock module failure, the system switches over to the redundant clock module automatically. This also results in a hardware reset of the switch. When the switch reboots, it displays the current active clock module. The following syslog message is printed at switch boot-up time, indicating the current active clock module.

Error Message `PLATFORM-0-CHASSIS_CLKSWRESET: Switch reset due to clock switch.`

Explanation Chassis clock source has failed and system will be reset. System will automatically start using the redundant clock module.

Recommended Action Replace the failed clock module during the next maintenance window.

Typically, clock module A is the active clock and on a failure of clock module A, clock module B becomes the active clock. Refer to the hardware installation guide for your platform at the following website to replace a clock module.

http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_installation_guides_list.html

Troubleshooting Other Hardware Issues

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

To identify a hardware issue with a module using the CLI, follow these steps:

Step 1 Use the **show module internal exceptionlog** CLI command.

The exception log is a wraparound log of all errors and exceptional conditions on each module. Some exceptions are catastrophic, some partially affect certain ports in a module, others are for warning purposes. Each log entry includes the following fields:

- device id—The device that logged the exception. This is interpreted by your customer support representative.
- device errorcode—The error code that occurred on the device. This is interpreted by your customer support representative.
- error type—The severity level of the error. Software errors are typically minor or warning. All other errors may be hardware problems.
- Number Ports went bad—The number of ports on the module that are no longer operational.
- system time— The timestamp when the problem occurred.

The exception log is stored in the NVRAM on the supervisor module.

Most hardware errors are logged in this command output. If the error type field indicates anything other than minor or warning error, then it is most likely a hardware failure. (See Example 3-5.)

Example 3-5 Output of show module internal exceptionlog Command

```
switch# show module internal exceptionlog
********* Exception info for module 6 ********

exception information --- exception instance 1 ----
```

```
device id:          85
device errorcode:   0xc550120c
system time:        (1127748710 ticks) Mon Sep 26 15:31:50 2005

error type:         Minor error
Number Ports went bad: none

********* Exception info for module 8 ******** <---Possible failed module

exception information --- exception instance 1 ----
device id:          12
device errorcode:   0x80000080
system time:        (1127843531 ticks) Tue Sep 27 17:52:11 2005

error type:         FATAL error <------------------- Error Type field
Number Ports went bad:
1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16
```

Step 2 View the error statistics from the **show hardware internal errors** command output.

Some error statistics reported under FC-MAC are not necessarily errors, but those counters normally do not increment for a port that is in an up state.

Step 3 View the interrupt counts in the **show hardware internal errors** command output.

Note the following:

- Some interrupts are not necessarily error interrupts.
- Some interrupts have a threshold before the corresponding ports are declared as bad. Do not conclude that the hardware is bad because of some interrupt counts. However, these commands are useful for your customer support representative when debugging the problems.
- Some interrupt counts may show up under UP-XBAR and DOWN-XBAR ASICs, when one of Supervisors is pulled out or restarted.

Troubleshooting Supervisor Issues

Supervisor initiation varies depending on whether or not you have a redundant supervisor present.When two supervisors are present in the system at poweredup, one of the supervisors will become active and the other standby. The active supervisor initialization differs from the standby supervisor.

If there is no active supervisor in the system, the supervisor that boots up first will default to the active supervisor. If there is an active supervisor in the system, the supervisor that is booting up will default to the standby supervisor state. The standby supervisor needs to mirror the state of the active supervisor. Once all the components on the standby are synchronized with that of the active supervisor, the standby supervisor is up.

Cisco SAN-OS maintains debug information during runtime. When a supervisor reboots, much of the debug information is lost. However, all critical information is stored in NVRAM and can be used to reconstruct the failure. When an active supervisor reboots, the information that is stored in its NVRAM cannot be obtained until it comes back up again. Once the supervisor reboots, use the following CLI commands can be used to view the persistent log:

- **show logging nvram**
- **show system reset-reason**
- **show module internal exception-log**

This section describes how to diagnose when an active or standby supervisor fails to initialize properly. This section includes the following topics:

Active Supervisor Reboots

Symptom Active supervisor reboots.

Table 3-7 Active Supervisor Reboots

Symptom	Possible Cause	Solution
Active supervisor reboots.	Supervisor process crashed, resulting in a supervisor reload.	Use the **show system reset-reason** CLI command to view the cause of the reset after the supervisor reboots. (See Example 3-6.) If you have a standby supervisor, the standby is now the active supervisor. Display the system message log on the standby supervisor to see the same information. (See Example 3-7.) Use the s**how process log** CLI command to view a list of process restarts.
	Runtime diagnostics failure detected.	Use the **show module internal exceptionlog** CLI command on the standby supervisor to view the cause of the reset after the supervisor reboots. (See Example 3-8.) If you have a standby supervisor, the standby is now the active supervisor. Display the system message log on the standby supervisor to see the same information. See (Example 3-9.) Optionally, when the supervisor reboots, use the **show system reset-reason** CLI command to view this same information.
		See also the "Troubleshooting Cisco SAN-OS Software System Reboots" section on page 2-12.

Example 3-6 displays the reset reason when a supervisor rebooted because of a process crash.

Example 3-6 Reset Reason for Supervisor Reboot Caused by Failed Process

```
switch# show system reset-reason
----- reset reason for module 6 -----
1) At 94009 usecs after Tue Sep 27 18:52:13 2005
    Reason: Reset triggered due to HA policy of Reset
    Service: Service "xbar" <------------------ Process that caused the reboot
    Version: 2.1(2)
```

Example 3-7 displays the system messages on the standby supervisor when a supervisor rebooted because of a process crash.

Example 3-7 System Messages for Supervisor Reboot Caused by Failed Process

```
Switch# show logging
2005 Sep 27 18:58:05 172.20.150.204 %SYSMGR-3-SERVICE_CRASHED: Service "xbar" (PID 1225)
hasn't caught signal 9 (no core).
2005 Sep 27 18:58:06 172.20.150.204 %SYSMGR-3-SERVICE_CRASHED: Service "xbar" (PID 2349)
hasn't caught signal 9 (no core).
2005 Sep 27 18:58:06 172.20.150.204 %SYSMGR-3-SERVICE_CRASHED: Service "xbar" (PID 2352)
hasn't caught signal 9 (no core).
```

Example 3-8 displays the exception log when a supervisor rebooted because of a runtime diagnostic failure.

Example 3-8 Exception Log for Supervisor Reboot Caused by Runtime Diagnostic Failure

```
switch# show module internal exceptionlog module 6
********* Exception info for module 6 ********

exception information --- exception instance 1 ----
device id:          12
device errorcode:   0x80000020
system time:        (1127917068 ticks) Wed Sep 28 14:17:48 2005

error type:         FATAL error <---------------------- exception that caused the reboot
Number Ports went bad:
1,2,3,4,5,6

exception information --- exception instance 2 ----
device id:          12
device errorcode:   0x00060a02
system time:        (1127917067 ticks) Wed Sep 28 14:17:47 2005

error type:         Warning
Number Ports went bad:
1,2,3,4,5,6
```

Example 3-9 displays the system messages on the standby supervisor when a supervisor rebooted because of a runtime diagnostic failure..

Example 3-9 System Messages for Supervisor Reboot Caused by Runtime Diagnostic Failure

```
Switch# show logging
2005 Sep 28 14:17:47 172.20.150.204 %XBAR-5-XBAR_STATUS_REPORT: Module 6 reported status
for component  12 code 0x60a02.
2005 Sep 28 14:17:59 172.20.150.204 %PORT-5-IF_UP: Interface mgmt0 on slot 5 is up
2005 Sep 28 14:18:00 172.20.150.204 %CALLHOME-2-EVENT: SUP_FAILURE
```

Standby Supervisor Not Recognized by Active Supervisor

Symptom Standby supervisor is not recognized by the active supervisor.

Table 3-8 Standby Supervisor Not Recognized by Active Supervisor

Symptom	Possible Cause	Solution
Standby supervisor not recognized by the active supervisor.	Standby supervisor did not sync properly with active supervisor.	See the "Standby Supervisor Not Recognized by Active Supervisor" section on page 3-18 to verify the problem. Observe the boot process to verify that the LEDs follow proper boot sequence and verify that the standby supervisor goes through the proper power-up, initializing, an testing phases. If the standby supervisor is at the `loader>` prompt, use the **reload module 6 force-dlnd** command from the active supervisor to force the standby supervisor to netboot off of the active supervisor.

Verifying That a Standby Supervisor Failed to Sync Using the CLI

To verify that a standby supervisor did not sync with the active supervisor using the CLI, follow these steps:

Step 1 Use the **show module** command on the active supervisor to verify that the active supervisor does not detect the standby supervisor.(See Example 3-10.)

Example 3-10 show module Command Output

```
switch#  show module
Mod  Ports  Module-Type                       Model              Status
---  -----  -------------------------------- ------------------ ------------
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
8    8      IP Storage Services Module                          powered-dn

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
5    2.1(2)       1.1     --


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
5    00-0b-be-f7-4d-1c to 00-0b-be-f7-4d-20  JAB070307XG

* this terminal session
```

Step 2 Telnet to the standby supervisor console port and verify that it is in standby mode. (See Example 3-11.)

Example 3-11 Verify Standby Supervisor Mode

```
runlog>telnet sw4-ts 2004
Trying 172.22.22.55...
Connected to sw4-ts.cisco.com (172.22.22.55).
Escape character is '^]'.
```

```
MDS Switch
 login: admin
Password:
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2005, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.
switch(standby)#
```

Step 3 Use the **show system redundancy status** CLI command on the active supervisor to verify that the standby supervisor did not complete the synchronization phase with the active supervisor.

```
switch#  show system redundancy status
Redundancy mode
---------------
      administrative:   HA
         operational:   None

This supervisor (sup-1)
-----------------------
    Redundancy state:   Active
    Supervisor state:   Active
      Internal state:   Active with HA standby

Other supervisor (sup-2)
------------------------
    Redundancy state:   Standby
    Supervisor state:   HA standby
      Internal state:   HA synchronization in progress
```

The most likely reason for the synchronization to stall is if one of the software components on the standby supervisor failed to synchronize its state with the active supervisor.

Step 4 Use the s**how system internal sysmgr gsyncstats** CLI command on the active supervisor to determine which processes did not synchronize on the standby supervisor.

```
switch# show system internal sysmgr gsyncstats
Name               Gsync done  Gsync time(sec)
----------------   ----------  -------------
aaa                     1          0
ExceptionLog            1          0
platform                1          1
radius                  1          0
securityd               1          0
SystemHealth            1          0
tacacs                  0        N/A
acl                     1          0
ascii-cfg               1          1
bios_daemon             0        N/A
bootvar                 1          0
callhome                1          0
capability              1          0
cdp                     1          0
cfs                     1          0
cimserver               1          0
cimxmlserver            0        N/A
confcheck               1          0
core-dmon               1          0
core-client             0        N/A
device-alias            1          0
```

```
dpvm                    0       N/A
dstats                  1         0
epld_upgrade            0       N/A
epp                     1         1
```

Step 5 Use the s**how system internal sysmgr service all** CLI command on the standby supervisor to determine if any process is experiencing excessive restarts. (See Example 3-12.)

Note This command may not be available if the standby supervisor is at the `loader>` prompt.

Example 3-12 Finding Excessive Restarts

```
switch(standby)# show system internal sysmgr service all
Name           UUID         PID    SAP    state   Start count
------------  --------     ------  -----   -----     -----------
aaa            0x000000B5    1458    111   s0009       1
ExceptionLog   0x00000050    [NA]   [NA]   s0002     None
platform       0x00000018    1064     39   s0009       1
radius         0x000000B7    1457    113   s0009       1
securityd      0x0000002A    1456     55   s0009       1
vsan           0x00000029    1436     15   s0009       1
vshd           0x00000028    1408     37   s0009       1
wwn            0x00000030     1435   114   s0009       1
xbar           0x00000017     [NA]   [NA]  s0017       23
xbar_client    0x00000049     1434   917   s0009        1
```

Looking at the standby supervisor in Example 3-12 shows that the crossbar (xbar) software component has been restarted 23 times. This has probably prevented the standby from initializing properly.

Step 6 Use the **reload module** CLI command to restart the standby supervisor. If this fails, use the **reload module 6 force-dlnd** command from the active supervisor to force the standby supervisor to netboot off of the active supervisor.

Standby Supervisor Stays in Powered-Up State

Symptom Standby supervisor stays in powered-up state.

Table 3-9

Symptom	Possible Cause	Solution
Standby supervisor stays in powered-up state.	Standby supervisor did not sync properly with active supervisor.	See the "Verifying That a Standby Supervisor Is in the Powered-Up State Using Device Manager" section on page 3-21 or the "Verifying That a Standby Supervisor Is in Powered-Up State Using the CLI" section on page 3-21.

Verifying That a Standby Supervisor Is in the Powered-Up State Using Device Manager

To verify that a standby supervisor is in the powered-up state using Device Manager, follow these steps:

Step 1 Choose **Physical > Modules....** and verify that the operational status of the standby supervisor (OperStatus) is PoweredUp.

Step 2 Right-click the standby supervisor and select **Reset** from the drop-down menu to restart the standby supervisor.

Verifying That a Standby Supervisor Is in Powered-Up State Using the CLI

To verify that a standby supervisor is in the powered-up state using the CLI, follow these steps:

Step 1 Use the **show module** command on the active supervisor to verify that the standby supervisor in the powered-up state.(See Example 3-13.)

Example 3-13 show module Command Output

```
switch# show module
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1                                  powered-up
8    8      IP Storage Services Module                           powered-dn

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
5    2.1(2)       1.1     --


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
5    00-0b-be-f7-4d-1c to 00-0b-be-f7-4d-20  JAB070307XG

* this terminal session
```

Step 2 Use the **show module internal event-history module** CLI command to determine what component may have failed.

Step 3 Use the **reload module** CLI command to restart the standby supervisor.

Troubleshooting Supervisor Modules

Note If only one supervisor module is installed, ensure that automatic synchronization is off before servicing the other module. This prevents the switch from attempting to fail over to an unavailable module.

This section provides a workaround for a failed supervisor under certain conditions. An example situation is used to describe the problem and the workaround.

In this case, the supervisor failed when the standby was reloaded, or when the supervisor was replaced with a new one. It was discovered that the failed supervisor either had its version of code changed, or the running configuration on the active supervisor was not saved with the appropriate boot parameters. In either case, the problem was mismatched code on the active and standby supervisors. One clue that indicated the mismatched code was a heartbeat error on the active supervisor. Because of this error, the current Flash images were unable to be copied from the active supervisor to the standby.

The workaround was to copy the images to CompactFlash, switch consoles, and load code from CompactFlash onto the second supervisor. The second supervisor was at a loader prompt, which is indicative of missing boot statements. When a **dir slot0:** CLI command was executed, none of the images appeared. This may have been the result of mismatched images on supervisors or to not having current images in Flash memory on the supervisor. Performing a **copy slot0: bootflash:** CLI command copied the images anyway. Once the images were loaded on the second supervisor and the boot statements were confirmed and saved on the active supervisor, the supervisor loaded and came up in standby-ha mode.

Troubleshooting Switching and Services Modules

This section describes problems with switching and services modules and includes the following topics and symptoms:

Overview of Module Status

Choose **Physical > Modules...** on Device Manager or use the **show module** CLI command to see the status of any module in a switch. (See Example 3-14.)

Example 3-14 show module Command Output

```
switch# show module 8
Mod  Ports  Module-Type                       Model              Status
---  -----  --------------------------------- ------------------ ------------
8    8      IP Storage Services Module        DS-X9308-SMIP      ok

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
8    2.1(2)       0.206   21:c1:00:05:30:00:8f:5e to 21:c8:00:05:30:00:8f:5e


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
8      00-05-30-00-9e-fa to 00-05-30-00-9f-06  JAB064704LH
```

The module status gives a good indication of the state of the module. Table 3-10 identifies all the different states that a module can experience and a brief description of the state.

Table 3-10 Module States

Module Status	Description	Module Status Condition
OK	The module is up and running.	Good
powered-down	The module has been powered down because of user configuration or an error. Us e the **show running-config | include poweroff** CLI command to determine if the module has been configured as powered-down. Otherwise, the module was powered down because of a failure. If a module reports a FATAL error, the supervisor logs an exception and reboots the module. If the supervisor reboots the module for errors three times in a one-hour interval, the supervisor keeps the module permanently powered down.	Good
err-pwd-dn		Failed
pwr-denied	The chassis does not have enough remaining power to power up the module. Use the **show environment power** CLI command to show the current power status of the switch.	Failed
powered-up	The module powered up and the supervisor is waiting for the module to initialize.	Transient
pwr-cycled	The module reloaded.	Transient
testing	The module has powered up and doing runtime diagnostics.	Transient
initializing	The module is receiving configuration from the supervisor.	Transient
upgrading	The module is in the process of a nondisruptive upgrade.	Transient
failure	The module has experienced a failure, but the module has not been power cycled because the debug flag was configured. Use the debug flag to collect debug information from the module as required by your customer support representative. Once all necessary data is collected, reload the module by using the **reload module** CLI command.	Failed

Module Initialization Overview

When a module is inserted into the switch, the module goes through an initial bring up sequence. This sequence bring the module to a known good state before the module is declared online. The initialization sequence includes the following steps:

- Module Bootup, page 3-24
- Image Download, page 3-24
- Runtime Diagnostics, page 3-25
- Runtime Configuration, page 3-25
- Online and Operational, page 3-25

Most of the module related failures (such as module not coming up, module getting reloaded, and so on) can be analyzed by looking at the logs stored on the switch. Use the following CLI commands to view this information:

- **show version**

- **show logging**
- **show module internal exception-log**
- **show module internal event-history module**
- **show module internal event-history errors**
- **show platform internal event-history errors**
- **show platform internal event-history module**

Module Bootup

When a module is inserted into the switch, the supervisor puts the module in powered-up state. In this state, the supervisor waits for the module to bootup and send its identification to the active supervisor.

If the supervisor does not receive the registration from the module within a given time frame, it power cycles the module. This failure is called a boot-up failure. The failure codes for boot-up failure can be obtained using the **show platform internal event-history errors** CLI command. (See Example 3-15.)

Example 3-15 Finding Boot-Up Failure Codes

```
switch# show platform internal event-history errors
The following error codes are defined
No Boot Device = 0xF1
Boot Failed= 0xC0
Net Boot Failed   = 0xD0
Unknown Status   = 0x1B
```

Image Download

Once the supervisor receives the registration message, it checks the image compatibility matrix. The image compatibility determines whether the version of code running on the supervisor is compatible with the version of code running on the module. If they do not match, the module downloads an updated version of the code, reboots, and sends a registration message again with the updated parameters.

If the module is unable to download the code, the supervisor generates the following system message:

Error Message `MODULE-2-MOD_DNLD_FAIL: Image download failed for module [dec].`

Explanation The module failed to download a new image from the supervisor module.

Recommended Action Collect module information by entering the **show module internal all module** <dec> command.

In addition the module generates a system message indicating the exact reason why the image download failed:

Error Message `IMAGE_DNLD-SLOT#-2-ADDON_IMG_DNLD_FAILED: Module image download process failed. [chars].`

Explanation The add-on image download to the module failed. This module is not operational until an add-on image has been successfully installed.

Recommended Action Verify the location and version of your module image. Enter **install module** CLI command or similar Fabric Manager/Device Manager command to download a new module image.

If the image download fails, the supervisor power cycles the module. Choose **Logs > Switch Resident > Syslog > Since Reboot** in Device Manager or use the **show logging** CLI command to view the failure messages.

Runtime Diagnostics

After the module succeeds registering with the supervisor, the module checks the hardware. If this fails, the module reports the error to the supervisor and generates the following system message:

Error Message `MODULE-2-MOD_DIAG_FAIL: Module [dec] reported failure on ports [dec]/[dec]-[dec]/[dec] ([chars]) due to [chars] in device [dec] (device error [hex]).`

Explanation The module reported a failure in the runtime diagnostic. Module manager is going to power cycle the module.

Recommended Action Collect information about the module by entering the **show module internal all module** CLI command.

In addition, this information is stored in the exception log (which is persistent across reboots). The supervisor then power cycles the module. Choose **Logs > Switch Resident > Syslog > Since Reboot** in Device Manager or use the **show logging** and **show module internal exception-log module** CLI commands to retrieve failure information.

Runtime Configuration

After the runtime diagnostics complete successfully, the module informs the supervisor that it is ready for configuration. Individual supervisor components configure the module. If any component reports a problem during this stage, the supervisor reboots the module. Use the **show module internal event-history module** CLI command to determine which component reported the problem.

Online and Operational

Once all the supervisor components have configured the module, the module goes to the ok state. In this state, the module is online and operational. The supervisor continues to monitor the module periodically to verify correct operation. The following events are monitored:

- Heartbeat message—Sent between the supervisor and the module to verify that the module is running.

- Online health management (OHMS)— Sent from the supervisor to all the ports in the module to verify that traffic is flowing properly.

In addition, the module monitors itself and generates an exception if it detects an anomalous condition. If the exception is a FATAL error, the module is power cycled. Use the following CLI commands to view the conditions leading up to the problem:

- **show logging**
- **show module diag**
- **show module internal exception-log module**
- **show module internal event-history module**
- **show hardware internal errors**

Analyzing The Logs

In some instances, you may need to check other internal logs to verify the cause of a problem. You can use the state transition log and the error log in these instances. These logs may hold information not present in the system messages or exception log because of interactions between the module and the supervisor. The state transition log is sorted in ascending manner (i.e. the latest state is at the end of the log). The error log is sorted in descending manner (that is, the latest error is at the beginning of the log).

Use the **show module internal event-history module** CLI command to view the state transition log for a module. Use the **show module internal event-history errors** CLI command to view the error log.

The state transition log indicates the current state of a given module. (See Example 3-16.) Each element of the transition log contains the following information:

- Timestamp
- Node that triggered the state transition
- Module state prior to transition
- Event that occurred
- Current state of module

Example 3-16 State Transition Log

```
7) FSM:<ID(2): Slot 8, node 0x0800> Transition at 14258 usecs after Mon Sep 26 17:50:56
2005
    Previous state: [LCM_ST_LC_POWERED_UP]
    Triggered event: [LCM_EV_PFM_LC_STATUS_POWERED_DOWN]
    Next state: [LCM_ST_LC_NOT_PRESENT]
```

Based on the above state transition you can infer that when the module was in the *powered-up* state, an event from PFM to power down the module was triggered. This trigger caused the state machine to go to the *not present* state.

Troubleshooting Module Issues

To isolate a module problem, follow these steps:

Step 1 Verify that all Status LEDs are green. If any status LED is red or off, the module might have shifted out of its slot.

Step 2 Reseat the module until both ejector levers are at 90 degrees to the rear of the chassis.

Step 3 Tighten the captive screws at the left and right of the module front panel.

Step 4 Restart the system.

If the Status LED on a switching module is orange, the module might be busy or disabled. Refer to the following website for the latest Cisco MDS 9000 Family configuration guides to configure or enable the interfaces:
http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm.
After the system reinitializes the interfaces, the Status LED on the module should be green.

Step 5 If the module does not transition into the online state, see the symptoms listed in this section.

If you are unable to resolve a problem with the startup, gather the information listed under Appendix A, "Before Contacting Technical Support" and contact your technical support representative for assistance as directed in the "Obtaining Technical Assistance" section on page xxv.

Troubleshooting Powered-Down Modules

Symptom Module is in the powered-down state.

The following system messages may be present if a module fails to power up:

Error Message `PLATFORM-2-PFM_LC_BOOT_DEV_ABSENT: No bootflash found in Module [dec].`

Explanation No bootflash found.

Recommended Action Put bootflash in the module and try again.

Error Message `PLATFORM-2-PFM_LC_BOOT_DEV_FAIL: BAD Bootflash found in Module [dec].`

Explanation Bad bootflash found.

Recommended Action Replace the bootflash in the module and try again.

Error Message `PLATFORM-2-PFM_LC_NETBOOT_FAIL: Netboot for Module [dec] failed.`

Explanation Netboot failed.

Recommended Action Replace the BIOS in the module. See the "Troubleshooting Cisco SAN-OS Software System Reboots" section on page 2-12.

Error Message `PLATFORM-2-PFM_LC_REGISTRATION_FAIL: Could not register with Module [dec].`

Explanation Module registration failed.

Recommended Action Replace the module.

Error Message `PLATFORM-2-PFM_LC_STATUS: Module [dec] powered up with [dec] status.`

Explanation Status for module that failed registration.

Recommended Action Replace the module.

Error Message `PLATFORM-3-MOD_PWRFAIL: Module [dec] failed to power up (Serial No. [chars]).`

Explanation The module failed to power up.

Recommended Action Enter the **show platform internal all module** [dec] CLI command to collect more information.

Introduced Cisco MDS SAN-OS Release 1.2(2a).

Error Message `PLATFORM-3-MOD_PWRIDPROMFAIL: Module [dec] failed to power up due to idprom read error.`

Explanation The module cannot be powered up because of an IDPROM read error.

Recommended Action Enter the **show platform internal all module** [dec] and **show module internal all module** [dec] **show sprom module** [dec][dec] CLI command to read module IDPROM contents to collect more information.

Error Message `PLATFORM-5-MOD_PWRDN: Module [dec] powered down (Serial No. [chars]).`

Explanation The module is powered down.

Enter the **show module**, **show platform internal all module**[dec] and **show module internal all module** [dec] CLI command to collect more information if you suspect module has been powered down due to errors.

Table 3-11 Module is in the Powered-Down State

Symptom	Possible Cause	Solution
Module is in powered-down state.	Module experienced boot-up failures.	**Choose Logs > Switch Resident > Syslog > Sever Events** on Device Manager or use the **show logging** CLI command to verify bootup problems. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	Module failed to register with the supervisor.	Use the **show module internal event-history module** CLI command and look for: `Triggered event: [LCM_EV_LCP_REGISTRATION_TIMEOUT]` to verify that the module did not register. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	Module failed to connect to fabric.	Use the **show system internal xbar internal event-history module** CLI command and look for : `Triggered event: [XBM_MOD_EV_SYNC_FAILED]` to verify that the module could not connect to the fabric. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	Supervisor failed to configure the module.	Verify the cause of the failure. See the "Diagnosing a Powered-Down Module" section on page 3-29. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.

Diagnosing a Powered-Down Module

To diagnose the reason for a powered-down module, follow these steps:

Step 1 Use the **show module** CLI command to verify the status of the module.

```
switch# show module
Mod  Ports  Module-Type                       Model                 Status
---  -----  --------------------------------- ------------------ ------------
5    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    ha-standby
6    0      Supervisor/Fabric-1               DS-X9530-SF1-K9    active *
8    8      IP Storage Services Module                           powered-dn

Mod  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  ---------------------------------

Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
```

```
5    00-0b-be-f7-4d-1c to 00-0b-be-f7-4d-20  JAB070307XG
6    00-05-30-00-93-7e to 00-05-30-00-93-82  JAB0637059v
```

Step 2 Use the **show logging** CLI command to see what events occurred on this module.

```
Switch# show logging

2005 Sep 27 15:26:02 172.20.150.204 %PLATFORM-5-MOD_DETECT: Module 8 detected (Serial
number JAB064704LH)
2005 Sep 27 15:26:02 172.20.150.204 %PLATFORM-5-MOD_PWRUP: Module 8 powered up (
Serial number JAB064704LH)
2005 Sep 27 15:27:03 172.20.150.204 %MODULE-5-MOD_REINIT: Re-initializing module 8
2005 Sep 27 15:27:09 172.20.150.204 %PLATFORM-5-MOD_DETECT: Module 8 detected (Serial
number JAB064704LH)
2005 Sep 27 15:27:09 172.20.150.204 %PLATFORM-5-MOD_PWRUP: Module 8 powered up (
Serial number JAB064704LH)
2005 Sep 27 15:28:10 172.20.150.204 %MODULE-5-MOD_REINIT: Re-initializing module 8
2005 Sep 27 15:28:15 172.20.150.204 %PLATFORM-5-MOD_DETECT: Module 8 detected (Serial
number JAB064704LH)
2005 Sep 27 15:28:15 172.20.150.204 %PLATFORM-5-MOD_PWRUP: Module 8 powered up (
Serial number JAB064704LH)
2005 Sep 27 15:29:16 172.20.150.204 %MODULE-5-MOD_REINIT: Re-initializing module 8
2005 Sep 27 15:29:22 172.20.150.204 %PLATFORM-5-MOD_DETECT: Module 8 detected (Serial
number JAB064704LH)
```

Note that module 8 powered up and reinitialized three times. This indicates that the module was never able to go online. The supervisor powered down the module.

Step 3 Use the **show module internal exception module** CLI command to view the exception log.

```
switch# show module internal exceptionlog module 8
********* Exception info for module 8 ********

exception information --- exception instance 1 ----
device id:          8
device errorcode:   0x40000002
system time:        (1127835023 ticks) Tue Sep 27 15:30:23 2005

error type:         Warning
Number Ports went bad: none

exception information --- exception instance 2 ----
device id:          8
device errorcode:   0x40000002
system time:        (1127834956 ticks) Tue Sep 27 15:29:16 2005

error type:         Warning
Number Ports went bad: none

exception information --- exception instance 3 ----
device id:          8
device errorcode:   0x40000002
system time:        (1127834890 ticks) Tue Sep 27 15:28:10 2005

error type:         Warning
Number Ports went bad: none

exception information --- exception instance 4 ----
device id:          8
device errorcode:   0x40000002
system time:        (1127834823 ticks) Tue Sep 27 15:27:03 2005
```

Note that the time when the module was reinitialized (from system messages) and the time when the exceptions were raised (in the exception log) are correlated. This means that device ID:8 had errors while bringing the module up.

Step 4 Use the **show module internal event-history module** CLI command to gather more information.

```
Switch# show module internal event-history module 8
79) Event:ESQ_START length:32, at 665931 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x2710, Ret:success
Seq Type:SERIAL

80) Event:ESQ_REQ length:32, at 667362 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x1, Ret:success
[E_MTS_TX] Dst:MTS_SAP_ILC_HELPER(125), Opc:MTS_OPC_LC_IS_MODULE_SAME(2810)

81) Event:ESQ_REQ length:32, at 667643 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x2, Ret:success
[E_MTS_TX] Dst:MTS_SAP_MIGUTILS_DAEMON(949), Opc:MTS_OPC_LC_INSERTED(1081)

82) Event:ESQ_RSP length:32, at 673004 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x2, Ret:success
[E_MTS_RX] Src:MTS_SAP_MIGUTILS_DAEMON(949), Opc:MTS_OPC_LC_INSERTED(1081)

83) Event:ESQ_REQ length:32, at 673265 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x3, Ret:success
[E_MTS_TX] Dst:MTS_SAP_XBAR_MANAGER(48), Opc:MTS_OPC_LC_INSERTED(1081)

85) Event:ESQ_RSP length:32, at 692394 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x3, Ret:(null)
[E_MTS_RX] Src:MTS_SAP_XBAR_MANAGER(48), Opc:MTS_OPC_LC_INSERTED(1081)

86) FSM:<ID(3): Slot 8, node 0x0802> Transition at 692410 usecs after Tue Sep 27
15:30:23 2005
Previous state: [LCM_ST_CHECK_INSERT_SEQUENCE]
Triggered event: [LCM_EV_LC_INSERTED_SEQ_FAILED]
Next state: [LCM_ST_CHECK_REMOVAL_SEQUENCE]

87) Event:ESQ_START length:32, at 692688 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x2710, Ret:success
Seq Type:SERIAL

88) Event:ESQ_REQ length:32, at 696483 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x1, Ret:success
[E_MTS_TX] Dst:MTS_SAP_MIGUTILS_DAEMON(949), Opc:MTS_OPC_LC_REMOVED(1082)

89) Event:ESQ_RSP length:32, at 698390 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0x1, Ret:success
[E_MTS_RX] Src:MTS_SAP_MIGUTILS_DAEMON(949), Opc:MTS_OPC_LC_REMOVED(1082)

108) Event:ESQ_REQ length:32, at 715171 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0xc, Ret:success
[E_MTS_TX] Dst:MTS_SAP_XBAR_MANAGER(48), Opc:MTS_OPC_LC_REMOVED(1082)

109) Event:ESQ_RSP length:32, at 716623 usecs after Tue Sep 27 15:30:23 2005
Instance:3, Seq Id:0xc, Ret:success
[E_MTS_RX] Src:MTS_SAP_XBAR_MANAGER(48), Opc:MTS_OPC_LC_REMOVED(1082)

110) FSM:<ID(3): Slot 8, node 0x0802> Transition at 716643 usecs after Tue Sep 2
7 15:30:23 2005
Previous state: [LCM_ST_CHECK_REMOVAL_SEQUENCE]
Triggered event: [LCM_EV_ALL_LC_REMOVED_RESP_RECEIVED]
Next state: [LCM_ST_LC_FAILURE]

111) FSM:<ID(3): Slot 8, node 0x0802> Transition at 716886 usecs after Tue Sep 2
```

```
7 15:30:23 2005
Previous state: [LCM_ST_LC_FAILURE]
Triggered event: [LCM_EV_LC_INSERTED_SEQ_FAILED]
Next state: [LCM_ST_LC_FAILURE]

112) FSM:<ID(3): Slot 8, node 0x0802> Transition at 717250 usecs after Tue Sep 2
7 15:30:23 2005
Previous state: [LCM_ST_LC_FAILURE]
Triggered event: [LCM_EV_FAILED_MORE3TIMES]
Next state: [LCM_ST_LC_NOT_PRESENT]

113) FSM:<ID(3): Slot 8, node 0x0802> Transition at 21633 usecs after Tue Sep 27
15:30:24 2005
Previous state: [LCM_ST_LC_NOT_PRESENT]
Triggered event: [LCM_EV_MODULE_POWERED_DOWN]
Next state: [LCM_ST_LC_NOT_PRESENT]


Curr state: [LCM_ST_LC_NOT_PRESENT]
```

Step 5 Starting with the most recent time (end of the log) and moving backwards in this example, you can infer the following:

```
Curr state: [LCM_ST_LC_NOT_PRESENT]<---- Indicates that the module is not present.

Index 112) Triggered event: [LCM_EV_FAILED_MORE3TIMES] <----Indicates that the module
failed repeatedly.

Index 111) Triggered event: [LCM_EV_LC_INSERTED_SEQ_FAILED] <---Indicates that the
insertion sequence failed.

Index 86) Previous state: [LCM_ST_CHECK_INSERT_SEQUENCE]
Triggered event: [LCM_EV_LC_INSERTED_SEQ_FAILED]
Next state: [LCM_ST_CHECK_REMOVAL_SEQUENCE] <---- Indicate that when module was being
inserted, the insertion failed and the module was removed.

Index 85)  Event:ESQ_RSP length:32, at 692394 usecs after Tue Sep 27 15:30:23 2005
   Instance:3, Seq Id:0x3, Ret:(null)
[E_MTS_RX] Src:MTS_SAP_XBAR_MANAGER(48),
Opc:MTS_OPC_LC_INSERTED(1081) <---Indicates the  event that caused the  module insertion
to fail. This indicates that xbar_manager failed.
```

In this example, you can conclude that module is not coming up, because the XBAR Manager is failing during the insertion of the module.

Troubleshooting Reloaded Modules

Symptom Module is automatically reloaded.

The following system messages may be present if a module reloads:

Error Message `MODULE-2-MOD_NOT_ALIVE: Module [dec] not responding... resetting.`

Explanation The module is not replying to the hello message. The module manager will reset the module.

Recommended Action No action is required.

Error Message `MODULE-2-MOD_SOMEPORTS_FAILED: Module [dec] reported failure on ports [dec]/[dec]-[dec]/[dec] ([chars]) due to [chars] in device [dec] (error [hex]).`

Explanation Module reported a failure in the runtime diagnostic because of a failure in some of the ports.

Recommended Action Collect module information by entering the **show module internal all module** CLI command.

Error Message `MODULE-2-MOD_DIAG_FAIL: Module [dec] reported failure on ports [dec]/[dec]-[dec]/[dec] ([chars]) due to [chars] in device [dec] (device error [hex]).`

Explanation The module reported a failure in the runtime diagnostic. Module manager is going to power cycle the module.

Recommended Action Collect information about the module by entering the **show module internal all module** CLI command.

Error Message `SYSTEMHEALTH-2-OHMS_MOD_PORT_LB_TEST_FAILED: Module [dec] Port [dec] has failed loop back tests.`

Explanation Port loop-back test failure.

Recommended Action No action is required.

Error Message `SYSTEMHEALTH-2-OHMS_MOD_SNAKE_TEST_FAILED: Module [dec] has failed snake loopback tests.`

Explanation Snake test failure.

Recommended Action No action is required.

Table 3-12 Module is Automatically Reloaded

Symptom	Possible Cause	Solution
Module is automatically reloaded.	Module experienced heartbeat failures.	Choose **Logs > Switch Resident > Syslog > Sever Events** on Device Manager or use the **show logging** CLI command to verify bootup problems. Use the **show module internal event-history module** CLI command and look for `Triggered event: [LCM_EV_LCP_ALIVE_TIMEOUT]` to verify that the module did not respond to heartbeat requests. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	The module experienced runtime diagnostic failures.	Verify the cause of the failure. See the "Diagnosing a Reloaded Module" section on page 3-34. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	Module lost sync with the fabric.	Use the **show system internal xbar internal event-history errors** and look for something similar to: `Rx MTS_OPC_SSA_LOST_SYNC_SERIAL slot 8 fabric 0 link 0` to verify that the module lost sync with the fabric. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.

Diagnosing a Reloaded Module

To diagnose the reason for a reloaded module, follow these steps:

Step 1 Right-click the module and select **Module** on Device Manager or use the **show module** CLI command to verify the status of the module.

Step 2 Choose **Logs > Switch Resident > Syslog > Sever Events** on Device Manager or use the **show logging** CLI command to search for common reload problems.

Step 3 Use the **show module internal exception module** CLI command to view the exception log.

```
switch#  show module internal exceptionlog module 8
********* Exception info for module 8 ********
exception information --- exception instance 3 ----
device id:          0
device errorcode:   0x40730017
system time:        (1127843486 ticks) Tue Sep 27 17:51:26 2005

error type:         FATAL error
Number Ports went bad:
1,2,3,4,5,6,7,8

exception information --- exception instance 4 ----
device id:          5
device errorcode:   0x40730019
system time:        (1127843486 ticks) Tue Sep 27 17:51:26 2005

error type:         Minor error
Number Ports went bad:
8
```

Step 4 Use the **show module internal event-history module** CLI command to gather more information.

```
Switch# show module internal event-history module 8
84) FSM:<ID(3): Slot 8, node 0x0802> Transition at 755101 usecs after Tue Sep 27
 17:51:26 2005
    Previous state: [LCM_ST_LC_ONLINE]
    Triggered event: [LCM_EV_LCP_RUNTIME_DIAG_FAILURE]
    Next state: [LCM_ST_CHECK_REMOVAL_SEQUENCE]

85) Event:ESQ_START length:32, at 755279 usecs after Tue Sep 27 17:51:26 2005
    Instance:3, Seq Id:0x2710, Ret:success
    Seq Type:SERIAL
```

Troubleshooting Modules in an Unkown State

Symptom Module is in the unknown state.

Table 3-13 Module Is in an Unknown State

Symptom	Possible Cause	Solution
Module is in an unknown state.	Module experienced SPROM failures.	Verify the cause of the failure. See the "Diagnosing a Module in the Unknown State" section on page 3-36. Right-click on the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.

Diagnosing a Module in the Unknown State

To diagnose a module in the unknown state, follow these steps:

Step 1 Right-click the module and select **Module** on Device Manager or use the **show module** CLI command to verify the status of the module.

Step 2 Choose **Logs > Switch Resident > Syslog > Sever Events** on Device Manager or use the **show logging** CLI command to search for common problems.

Step 3 Use the **show platform internal event-history errors CLI command** to view possible causes for the unknown state.

```
switch# show platform internal event-history errors
1) Event:E_DEBUG, length:37, at 370073 usecs after Thu Sep 29 17:22:48 2005
    [103] unable to init lc sprom 0 mod 8

switch# show platform internal event-history module 8
Inside pfm_show_eventlog
Index 1 TOKEN ID: 927
Index 2 TOKEN ID: 910
Module number 0x8

>>>>FSM: <Slot 8> has 2 logged transitions<<<<<

1) FSM:<Slot 8> Transition at 500219 usecs after Thu Sep 29 17:22:43 2005
    Previous state: [PLTFRM_STATE_MODULE_ABSENT]
    Triggered event: [PLTFRM_EVENT_MODULE_INSERTED]
    Next state: [PLTFRM_STATE_MODULE_PRESENT]

2) FSM:<Slot 8> Transition at 370112 usecs after Thu Sep 29 17:22:48 2005
    Previous state: [PLTFRM_STATE_MODULE_PRESENT]
    Triggered event: [PLTFRM_EVENT_MODULE_BOOTUP_ERROR]
    Next state: [PLTFRM_STATE_MODULE_UNRECOVERABLE_ERROR]


    Curr state: [PLTFRM_STATE_MODULE_UNRECOVERABLE_ERROR]
```

Troubleshooting Modules Not Detected by the Supervisor

Symptom Module is not detected by the supervisor.

Table 3-14 Module Is Not Detected by Supervisor

Symptom	Possible Cause	Solution
Module is not detected by the supervisor.	Module experienced SPROM failures.	Verify the cause of the failure. Right-click the module in Device Manager and select **Reset** or use the **reload module** CLI command to restart the module. See the "Reinitializing a Failed Module Using Fabric Manager" section on page 3-38 or the "Reinitializing a Failed Module Using the CLI" section on page 3-39.
	Module is not supported by the current version of Cisco SAN-OS on the switch.	Upgrade the software version on the switch. See the "Installing SAN-OS Software Using Fabric Manager" section on page 2-9 or the "Installing Cisco SAN-OS Software from the CLI" section on page 2-10.

Diagnosing a Module Not Detected by the Supervisor

To diagnose a module that has not been detected by the supervisor, follow these steps:

Step 1 Right-click the module and select **Module** on Device Manager or use the **show module** CLI command to verify the status of the module.

Step 2 Choose **Logs > Switch Resident > Syslog > Server Events** on Device Manager or use the **show logging** CLI command to search for common problems.

Step 3 Use the **show platform internal event-history errors** CLI command to view possible causes.

```
switch# show platform internal event-history errors
1) Event:E_DEBUG, length:42, at 703984 usecs after Thu Sep 29 17:46:20 2005
    [103] Module 8 pwr mgmt I/O cntrl reg 0x74


2) Event:E_DEBUG, length:69, at 703888 usecs after Thu Sep 29 17:46:20 2005
    [103] Module 8 pwr mgmt rev reg 0x74 brd present but power ok not set

switch# show platform internal event-history module 8
Inside pfm_show_eventlog
Index 1 TOKEN ID: 927
Index 2 TOKEN ID: 910
Module number 0x8

>>>>FSM: <Slot 8> has 10 logged transitions<<<<<

1) FSM:<Slot 8> Transition at 370299 usecs after Thu Sep 29 17:46:12 2005
    Previous state: [PLTFRM_STATE_MODULE_ABSENT]
    Triggered event: [PLTFRM_EVENT_MODULE_INSERTED]
    Next state: [PLTFRM_STATE_MODULE_PRESENT]

2) FSM:<Slot 8> Transition at 698894 usecs after Thu Sep 29 17:46:17 2005
    Previous state: [PLTFRM_STATE_MODULE_PRESENT]
    Triggered event: [PLTFRM_EVENT_MODULE_SPROM_READ]
    Next state: [PLTFRM_STATE_MODULE_POWER_EVAL]

3) FSM:<Slot 8> Transition at 705551 usecs after Thu Sep 29 17:46:17 2005
```

```
    Previous state: [PLTFRM_STATE_MODULE_POWER_EVAL]
    Triggered event: [PLTFRM_EVENT_MOD_START_POWER_UP]
    Next state: [PLTFRM_STATE_MODULE_START_POWER_UP]

4) FSM:<Slot 8> Transition at 110120 usecs after Thu Sep 29 17:46:20 2005
    Previous state: [PLTFRM_STATE_MODULE_START_POWER_UP]
    Triggered event: [PLTFRM_EVENT_MOD_END_POWER_UP]
    Next state: [PLTFRM_STATE_MODULE_POWERED_UP]

5) FSM:<Slot 8> Transition at 704067 usecs after Thu Sep 29 17:46:20 2005
    Previous state: [PLTFRM_STATE_MODULE_POWERED_UP]
    Triggered event: [PLTFRM_EVENT_MODULE_REMOVED]
    Next state: [PLTFRM_STATE_MODULE_ABSENT]
```

When a module is inserted into the switch, the supervisor reads the SPROM contents of the module. If the module is supported by the current version of Cisco SAN-OS, the module will be powered-up by the supervisor. If the power status does not come up ok, the module information is not relayed to the supervisor.

Reinitializing a Failed Module Using Fabric Manager

To reinitialize a failed module using the Fabric Manager, follow these steps:

Step 1 Choose **Switches > Copy Configuration** to save the running configuration to the startup configuration.

Step 2 Choose **Switches > Hardware**. Then select the **Module Status** tab in the Information pane and check the **Reset** check box to reload the module. Click the **Apply Changes** icon.

Step 3 If the module is not up, choose **Switches > Hardware** and check the S/W Rev column to verify the software image on the module.

Step 4 If the software image on the module is not the latest, choose **Tools > Other > Software Install** to download the latest image to supervisor bootflash memory.

Step 5 Use the CLI to force-download the software image from the supervisor to the module.

```
switch# reload module 2 force-dnld
```

Step 6 If the module is still not up, choose **Switches > Hardware** and view the Power Admin column to verify the power status for the module.

Step 7 If the module is not powered on, remove and reseat the module and select **on** from the Power Admin drop-down menu to power on the module.

Step 8 If the module is still not up, right-click on the switch in the map pane and select **Reset** to reload the entire switch.

Reinitializing a Failed Module Using the CLI

To reinitialize a failed module using the CLI, follow these steps:

Step 1 Save the running configuration to the startup configuration.

```
switch# copy running-config start-config
```

Step 2 Reload the module.

```
switch# reload module 2
```

Step 3 If the module is not up, verify the software image on the module.

```
switch# show module
```

Step 4 If the software image on the module is not the latest, download the latest image to supervisor bootflash memory.

```
switch# copy tftp: bootflash:
```

Step 5 Force-download the software image from the supervisor to the module.

```
switch# reload module 2 force-dnld
```

Step 6 If the module is still not up, verify the power status for the module.

```
switch# show environment power
```

Step 7 If the module is not powered on, remove and reseat the module and then power on the module.

```
switch# config t
switch(config)# no poweroff module 2
switch(config)# exit
switch#
```

Step 8 If the module is still not up, reload the entire switch.

```
switch# reload
```

Module Resets

Resets and reboots of modules are covered in detail in the "Troubleshooting Cisco SAN-OS Software System Reboots" section on page 2-12. If you use the **module reset-reason** CLI command and the output has an "unknown" reset reason, this may indicate a hardware problem. Some of the conditions that may cause this include:

- The switch experienced a power reset. This may be because you reset the power supplies, or because of a power interruption or failure.
- The front panel reset button on the supervisor module was pressed.
- Any hardware failure that caused the processor, dynamic memory, or I/O to reset or hang.

CHAPTER 4

Troubleshooting Licensing

Licensing functionality is available in all switches in the Cisco MDS 9000 Family. This functionality allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are supported, and enforced in Cisco MDS SAN-OS Release 1.3(1) and later.

This chapter includes the following topics:

License Overview

Cisco SAN-OS requires licenses for advanced features. These licenses have two options:

- Feature-based licensing—Features that are applicable to the entire switch. You need to purchase and install a license for each switch that uses the features you are interested in. The Enterprise license is an example of a feature-based license.
- Module-based licensing—Features that require additional hardware modules. You need to purchase and install a license for each module that uses the features you are interested in. The SAN Extension over IP license is an example of a module-based license.

Note The Cisco MDS 9216i switch enables SAN Extension features on the two fixed IP services ports only. The features enabled on these ports are identical to the features enabled by the SAN Extension over IP license on the14/2-port Multiprotocol Services (MPS-14/2) module. If you install a module with IP ports in the empty slot on the Cisco MDS 9216i, a separate SAN Extension over IP license is required to enable related features on the IP ports of the additional module.

Chassis Serial Numbers

Licenses are created using the serial number of the chassis where the license file is to be installed. Once you order a license based on a chassis serial number, you cannot use this license on any other switch. If you use a license meant for another chassis, you may see the following system message:

Error Message `LICMGR-3-LOG_LIC_INVALID_HOSTID: Invalid license hostid VDH=[chars] for feature [chars].`

Explanation The feature has a license with an invalid license Host ID. This can happen if a supervisor module with licensed features for one switch is installed on another switch.

Recommended Action Reinstall the correct license for the chassis where the supervisor module is installed.

Grace Period

If you use a feature that requires a license but have not installed a license for that feature, you are given a 120 day grace period to evaluate the feature. You must purchase and install the number of licenses required for that feature before the grace period ends or Cisco SAN-OS will disable the feature at the end of the grace period. If you try to use an unlicensed feature, you may see the following system messages:

Error Message `LICMGR-2-LOG_LIC_GRACE_EXPIRED: Grace period expired for feature [chars].`

Explanation The unlicensed feature has exceeded its grace time period. Applications using this license will be shut down immediately.

Recommended Action Please install the license file to continue using the feature.

Error Message `LICMGR-3-LOG_LICAPP_NO_LIC: Application [chars] running without [chars] license, shutdown in [dec] days.`

Explanation The Application [chars1] has not been licensed. The application will work for a grace period of [dec] days after which it will be shut down unless a license file for the feature is installed.

Recommended Action Install the license to continue using the feature.

Error Message `LICMGR-3-LOG_LIC_LICENSE_EXPIRED: Evaluation license expired for feature [chars].`

Explanation The feature has exceeded its evaluation time period. The feature will be shut down after a grace period.

Recommended Action Install the license to continue using the feature.

Error Message `LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature [chars]. Application(s) shutdown in [dec] days.`

Explanation The feature has not been licensed. The feature will work for a grace period, after which the application(s) using the feature will be shutdown.

Recommended Action Install the license to continue using the feature.

Error Message `LICMGR-6-LOG_LICAPP_EXPIRY_WARNING: Application [chars] evaluation license [chars] expiry in [dec] days.`

Explanation The application will exceed its evaluation time period in the listed number of days and will be shut down unless a permanent license for the feature is installed.

Recommended Action Install the license file to continue using the feature.

License packages can contain several features.If you disable a feature during the grace period and there are other features in that license package that are still enabled, the clock does not stop for that license package. To suspend the grace period countdown for a licensed feature, you must disable every feature in that license package. Choose **Switches > Licenses** and select the **Usage** tab in Fabric Manager or use the **show license usage** CLI command to determine which features are enabled for a license package.

Best Practices

This section provides the best practices when dealing with licenses for Cisco SAN-OS products.

- Do not ignore grace period expiration warnings. Allow 60 days before the grace period expires to allow time for ordering, shipping, and installation.
- Carefully determine the license(s) you require based on the features and modules that require a license. Remember that you need one license per chassis for feature-based licenses and one per module for module-based licenses.
- Order your license accurately:
  - Enter the Product Authorization Key that appears in the Proof of Purchase document that comes with your switch.
  - Enter the correct chassis serial number when ordering the license. The serial number must be for the same chassis that you plan to install the license on. Choose **Switches > Hardware** and check the SerialNo Primary for the switch chassis in Fabric Manager or use the **show license host-id** CLI command.
  - Enter serial numbers accurately. The serial number contains zeros, but no letter “O".
  - Order the license specific to your chassis or module type. An MDS 9200 Series license will not work on an MDS 9500 Series switch. Similarly, the SAN_EXTENSION_OVER_IP2 license works for an MPS-14/2 module, but will not work for an IPS-4 module. See Table 8-3 on page 8-7 for details on the SAN Extension over IP licenses available.
- Install licenses using the one-click method in Fabric Manager.
- Backup the license file to a remote, secure place. Archiving your license files ensures that you will not lose the licenses in the case of a failure on your switch.
- Install the correct licenses on each switch, using the licenses that were ordered using that switch’s serial number. Licenses are serial-number specific and platform or module type specific.
- Choose **Switches > Licenses** and select the **Usage** tab in Fabric Manager or use the **show license usage** CLI command to verify the license installation.
- Never modify a license file or attempt to use it on a switch that it was not ordered for. If you RMA a chassis, contact your customer support representative to order a replacement license for the new chassis.

Initial Troubleshooting Checklist

Begin troubleshooting license issues by checking the following issues first:

Checklist	Checkoff
Verify the chassis serial number for all licenses ordered.	❑
Verify the platform or module type for all licenses ordered.	❑
Verify that the Product Authorization Key you used to order the licenses comes from the same chassis that you retrieved the chassis serial number on.	❑
Verify that you have installed all licenses on all switches that require the licenses for the features you enable.	❑

This section includes the following topics:

Displaying License Information Using Fabric Manager

To view license information using Fabric Manager, follow these steps:

Step 1 Select **Switches > Licenses** from the Physical Attributes pane. You see the license information in the Information pane, one line per feature.

Step 2 Click the **Feature Usage** tab to see the switch, name of the feature package, the type of license installed, the number of licenses used (Installed Count), the expiration date, the grace period (if you do not have a license for a particular feature), and any errors (for example, if you have a missing license). Click the **Keys** tab to display information about each of the License Key files installed on your switches.

Step 3 Click the **Usage** tab to see the applications using the feature package on each switch. Use this tab to determine which applications depend on each license you have installed.

Displaying License Information Using Fabric Manager Web Services

Fabric Manager Release 2.1(2) or later supports viewing license use across the fabric from Fabric Manager Web Services. This view summarizes the licenses used on all switches in the fabric.

To view licenses using Fabric Manager Web Services, choose **Inventory > Licenses**.

Displaying License Information Using the CLI

Use the **show license** commands to display all license information configured on this switch (see Example 4-1 through Example 4-3).

Example 4-1 Displays Information About Current License Usage

```
switch# show license usage
Feature                 Installed   License Status  ExpiryDate   Comments
                                    Count
--------------------------------------------------------------------------------
FM_SERVER_PKG           Yes         -       Unused   never       license missing
MAINFRAME_PKG           No          -       Unused               Grace Period 57days15hrs
ENTERPRISE_PKG          Yes         -       InUse    never       -
SAN_EXTN_OVER_IP        No          0       Unused               -
SAN_EXTN_OVER_IP_IPS4   No          0       Unused
--------------------------------------------------------------------------------
```

Example 4-2 Displays the List of Features in a Specified Package

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

Example 4-3 Displays the Host ID for the License

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

Note Use the entire ID that appears after the colon (:) . The VHD is the Vendor Host ID.

Example 4-4 Displays All Installed License Key Files and Contents

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
        HOSTID=VDH=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
Evaluation.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 30-Dec-2003 uncounted \
        HOSTID=VDH=FOX0646S017 \
        NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Example 4-5 Displays a List of Installed License Key Files

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

Example 4-6 Displays the Contents of a Specified License Key File

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

Licensing Installation Issues

Common problems with licenses usually stem from incorrectly ordering the license file, installing the license file on an incorrect switch, or not ordering the correct number of licenses for your fabric.

This section includes the following topics:

One-Click License Install Fails or Cannot Connect to Licensing Website

The one-click license installation tries to open an HTTPS connection to the licensing website that matches the vendor you purchased your switch from.

Symptom One-click license install fails or cannot connect to the licensing website.

Table 4-1 One-Click License Install Fails or Cannot Connect to License Website

Symptom	Possible Cause	Solution

One-click license install fails or cannot connect to the licensing website.	License website uses HTTP, not HTTPS.	Edit <*install directory*>/bin/FabricManager.bat file to add the following lines to the JVMargs argument:
		`-Dhttp.proxyHost=HOSTADDRESS`
		`-Dhttp.proxyPort=HOSTPORT.`

	Fabric Manager communicating through a proxy server.	Edit <*install directory*>/bin/FabricManager.bat file to add the following lines to the JVMargs argument:
		`-Dhttps.proxyHost=HOSTADDRESS`
		`-Dhttps.proxyPort=HOSTPORT.`

	Java versions 1.4.2_01 and later do not have the right set of Certificate Authority (CA) certificates to validate the SSL certificates on the EMC server (HTTPS).	The license wizard cannot make an HTTPS connection to the EMC servers. If the License Wizard fails to fetch the license keys, saying the connection failed, the workaround is to install the latest 1.4(x) version of Java, preferably 1.4.2_04 or later.

Serial Number Issues

A common problem with licenses stems from not using the correct chassis serial number when ordering your license.

To obtain the correct chassis serial number using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Hardware** and select the **Inventory** tab.

Step 2 Copy down the SerialNo Primary field for the chassis that matches where you want to install a new license.

Note If you are ordering a module-based license, such as the SAN Extension over IP license package, you still use the chassis serial number for the chassis where the module resides, not the module serial number.

Use the **show license host-id** CLI command to obtain the correct chassis serial number for your switch using the CLI.

When entering the chassis serial number during the license ordering process, do not use the letter “O” in place of any zeros in the serial number.

RMA Chassis Errors or License Transfers Between Switches

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

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

Receiving Grace Period Warnings After License Installation

If the license installation does not proceed correctly, or if you are using a feature that exists in a license package that you have not installed, you will continue to get grace period warnings.

Symptom Receiving grace period warnings after a license installation.

Table 4-2 Receiving Grace Period Warnings After License Installation

Symptom	Possible Cause	Solution
Receiving grace period warnings after a license installation.	License file copied to switch but not installed.	Choose **Tools > Other > License Install** in Fabric Manager or use the **license install** CLI command to install the license.
	License installation failed.	Check your logs for any system messages for a failed license installation. Choose **Switches > Licenses** and select the **Usage** tab in Fabric Manager or use the **show license usage** CLI command to determine what feature is in use without a license.
	Not enough license files installed for a feature.	Some features require more than one license per chassis. Module-based licenses such as SAN Extension over IP for example requires one license per module that uses these features. Choose **Switches > Licenses** and select the **Usage** tab in Fabric Manager or use the **show license usage** CLI command to determine which feature is in use without a license.

Incorrect Number of Licenses in Use for Multiple Modules

Module-based licenses require one license installed per module that uses a licensed feature. SAN Extension over IP is an example of a module based license. Installing a SAN Extension over IP license while two FCIP instances from different modules are present, may cause the system to return the following error message:

```
Installing license failed: Number of License in use is more than the number being
installed.
```

This error message is generated because the license grace period is only applicable when no licenses are installed. The installation of one license terminates the grace period and will arbitrarily cause the second module to shut down, because this is not allowed by licensing.

The workaround for this scenario includes doing one of the following:

- Concatenate both licenses into one license file.
- Manually reduce the usage count by one.

To concatenate both licenses into one license file, follow these steps:

Step 1 Open both license files using WordPad.

Step 2 Copy both license files to one file:

```
Example
SERVER this_host ANY
VENDOR cisco
INCREMENT SAN_EXTN_OVER_IP_IPS2 cisco 1.0 permanent 1 \
VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M9500EXT12EK9=</SKU> \
HOSTID=VDH=FOXYYYYYYY \
NOTICE="<LicFileID>2005082204514XXXX</LicFileID><LicLineID>1</LicLineID> \
<PAK>MDS-1X-JAB-0F1A81</PAK>" SIGN=F0652E02XXXX
INCREMENT SAN_EXTN_OVER_IP_IPS2 cisco 1.0 permanent 1 \
VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M9500EXT12EK9=</SKU> \
HOSTID=VDH=FOXYYYYYYY \
NOTICE="<LicFileID>2005082204572XXXX</LicFileID><LicLineID>1</LicLineID> \
<PAK>MDS-1X-JAB-0F1AD1</PAK>" SIGN=D222AE4AXXXX
```

Step 3 Save the new concatenated license file.

Step 4 Upload and install the concatenated license file on the MDS switch.

To reduce the usage count to one, follow these steps:

Step 1 Bring down one of the modules manually to reduce the usage count by one.

Step 2 Reinsert the module after installing both licenses.

Grace Period Alerts

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

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

The grace period operates across all features in a license package. License packages can contain several features. If you disable a feature during the grace period and there are other features in that license package that are still enabled, the countdown does not stop for that license package. To suspend the grace period countdown for a license package, you must disable every feature in that license package. To disable the grace period countdown for Fabric Manager Server, you must explicitly check in the license using Device Manager. See the "Checking in the Fabric Manager Server License From Device Manager" section on page 4-10.

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

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

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

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

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

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

Checking in the Fabric Manager Server License From Device Manager

If you evaluated Fabric Manager Server without a license, you can stop the grace period countdown and disable all features using the Fabric Manager Server license package using Device Manager.

To stop the Fabric Manager Server license grace period using Device Manager, follow these steps:

Step 1 Choose **Admin > Licenses** and select the **Features** tab.

Step 2 Click **Check In FM**.

Note This button appears only when FM_SERVER_PKG is unlicensed.

Note Because of Caveat CSCeg23889, you might still receive Call Home or system messages for an unused FM_SERVER_PKG license. This caveat describes how extraneous messages are sent after a Fabric Manager Server license is checked in.

License Listed as Missing

After a license is installed and operating properly, it may show up as missing if you modify your system hardware or encounter a bootflash: issue.

Symptom License listed as missing.

Table 4-3 License Listed as Missing

Symptom	Possible Causes	Solutions
License listed as missing.	Supervisor module was replaced after license was installed.	Reinstall the license.
	Supervisor bootflash: is corrupted.	See the "Corrupted Bootflash Recovery" section on page 2-13 to recover from corrupted bootflash:. Reinstall the license.

CHAPTER 5

Troubleshooting Cisco Fabric Services

This chapter describes procedures used to troubleshoot Cisco Fabric Services (CFS) problems in the Cisco MDS 9000 Family multilayer directors and fabric switches. It includes the following sections:

Overview

Many features in the Cisco MDS 9000 Family switches require configuration synchronization in all switches in the fabric. It is important to maintain configuration synchronization across a fabric for consistency. In the absence of a common infrastructure, such synchronization is achieved through manual configuration at each switch in the fabric. This process is tedious and error prone.

As of Cisco MDS SAN-OS Release 2.0(1b), Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the applications. CFS can discover CFS-capable switches in the fabric as well as their application capabilities. Applications that can be synchronized using CFS include:

- IVR
- NTP
- DPVM
- user roles
- AAA server addresses
- syslog
- call home

Applications may be add to this list in future releases.

All switches in the fabric must be CFS capable. A Cisco MDS 9000 Family switch is CFS capable if it is running Cisco SAN-OS Release 2.0(1b) or later. Switches that are not CFS capable do not receive distributions and result in part of the fabric not receiving the intended distribution.

CFS has the following requirements:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the fabric.
- Pending database—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the pending database overwrites the configuration database (also know as active database or the effective database).
- CFS distribution enabled or disabled on a per-application basis—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the fabric.
- Explicit CFS commit—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database and distributes the new database to the fabric and releases the fabric lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.

As of Cisco SAN-OS Release 2.1(2b), some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.

Best Practices

You can avoid problems when configuring CFS if you observe the following best practices:

- Make sure that all the applications that you are using are enabled for CFS distribution on all the switches. By doing so, you ensure that application specific configurations will be in sync across the fabric.
- Do not simultaneously acquire a lock by configuring CFS from two different switches for the same application, even though the CFS module is capable of handling this type of activity. Applications on both sides might try to take the lock and might take a while to come out of the deadlock.
- If the CFS distribution for an application is enabled, then ensure that you either commit, abort, or clear the changes once you start the configuration. Applications take the lock on all the switches that come under the scope of the application's distribution. Once the lock is taken, if there is an ISL flap or a new switch joins the fabric, then the merge for that application goes into the waiting/in progress state until the lock is released.

Initial Troubleshooting Checklist

Begin troubleshooting CFS issues by checking the following issues first:

Checklist	Checkoff
Verify that CFS is enabled for the same applications on all affected switches.	❑
Verify that CFS distribution is enabled for the same applications on all affected switches.	❑
Verify that there are no pending changes for an application and that a CFS commit was issued for any configuration changes in a CFS enabled application.	❑
Verify that there are no unexpected CFS locked sessions. Clear any unexpected locked sessions.	❑

This section includes the following topics:

Verifying CFS Using Fabric Manager

To verify CFS using Fabric Manager or Device Manager, follow these steps:

Step 1 Choose **Admin > CFS** on Device Manager to verify that an application is listed and enabled. Repeat this on all switches.

Step 2 To list the set of switches in which an application is registered with CFS, choose the application configuration menu on Fabric Manager and select the CFS tab. For example, to verify that DPVM is enabled and global distribution is enabled on all switches, choose **Fabricxx > All VSANs > DPVM** and select the **CFS** tab. Verify that the Oper field is enabled and the Global filed is enabled for all switches in the fabric.

Step 3 To determine if all the switches in the fabric constitute one CFS fabric, or a multitude of partitioned CFS fabrics using Device Manager, follow these steps:

a. Choose **Admin > CFS** and highlight the application that you want to verify CFS on.

b. Click **Details** and select the **Merge** tab in the Details dialog box.

c. If you see multiple rows in the Merge status table, then the fabric is partitioned into multiple CFS fabrics. Some features enable CFS per VSAN and this is expected. If the selected feature should be fabric wide but you see multiple rows in the Merge status table, then the fabric may be partitioned , and the merge status may show that the merge has failed, is pending, or is waiting.

Verifying CFS Using the CLI

To verify CFS using the CLI, follow these steps:

Step 1 To verify that an application is listed and enabled, issue the **show cfs application** command to all switches. An example of the **show cfs application** command follows:

```
Switch# show cfs application

----------------------------------------
 Application    Enabled   Scope
----------------------------------------
ivr             Yes       Physical
ntp             No        Physical
dpvm            Yes       Physical
fscm            Yes       Physical
role            Yes       Physical
radius          Yes       Physical
fctimer         No        Physical
syslogd         No        Physical
callhome        No        Physical
device-alias    Yes       Physical
port-security   Yes       Logical

Total number of entries = 11
```

The Physical scope means that CFS applies the configuration for that application to the entire switch. The Logical scope means that CFS applies the configuration for that application to a specific VSAN.

Step 2 Verify the set of switches in which an application is registered with CFS, using the **show cfs peers name** *application-name* for physical scope applications, and the **show cfs peers name** *application-name* **vsan** *vsan-id* for logical scope applications.

An example command output for a physical scope application follows:

```
Switch# show cfs peers name dpvm

Scope      : Physical
--------------------------------------------------
 Switch WWN               IP Address
--------------------------------------------------
 20:00:00:0e:d7:0e:bf:c0  10.76.100.51    [Local]
 20:00:00:0e:d7:00:3c:9e  10.76.100.52

Total number of entries = 2
```

Note The **show cfs peers name** *application-name* command displays the peers for all VSANs when applied to a logical application.

An example command output for a logical scope application follows:

```
Switch# show cfs peers name port-security

Scope      :Logical [VSAN 1]
------------------------------------------------------------
 Domain   Switch WWN               IP Address
------------------------------------------------------------
 236      20:00:00:0e:d7:00:3c:9e  10.76.100.52    [Local]
 239      20:00:00:05:30:00:6b:9e  10.76.100.167
 101      20:00:00:0d:ec:06:55:c0  10.76.100.205
```

```
Total number of entries = 3


Scope      :Logical [VSAN 2]
-----------------------------------------------------------
 Domain   Switch WWN               IP Address
-----------------------------------------------------------
 239      20:00:00:0e:d7:00:3c:9e  10.76.100.52    [Local]
 211      20:00:00:05:30:00:6b:9e  10.76.100.167
 110      20:00:00:0d:ec:06:55:c0  10.76.100.205

Total number of entries = 3

Scope      :Logical [VSAN 3]
-----------------------------------------------------------
 Domain   Switch WWN               IP Address
-----------------------------------------------------------
 103      20:00:00:0e:d7:00:3c:9e  10.76.100.52    [Local]
 221      20:00:00:05:30:00:6b:9e  10.76.100.167
 11       20:00:00:0d:ec:06:55:c0  10.76.100.205

Total number of entries = 3
```

Step 3 To determine if all the switches in the fabric constitute one CFS fabric, or a multitude of partitioned CFS fabrics, issue the **show cfs merge status name** *application-name* command and the **show cfs peers name** *application-name* command and compare the outputs. If the outputs contain the same list of switches, the entire set of switches constitutes one CFS fabric. When this is the case the merge status should always show success at all switches. Example command outputs follow:

```
Switch# show cfs merge status name dpvm
Physical  Merge Status: Success [ Sat Nov 20 11:59:36 2004 ]
 Local Fabric
-----------------------------------------------------------
 Switch WWN               IP Address
-----------------------------------------------------------
 20:00:00:05:30:00:4a:de  10.76.100.51    [Merge Master]
 20:00:00:0d:ec:0c:f1:40  10.76.100.204


Switch# show cfs peers name dpvm

Scope      : Physical
-------------------------------------------------
 Switch WWN               IP Address
-------------------------------------------------
 20:00:00:0d:ec:0c:f1:40  10.76.100.204   [Local]
 20:00:00:05:30:00:4a:de  10.76.100.51

Total number of entries = 2
```

If the list of switches in the **show cfs merge status name** command output is shorter than that of the **show cfs peers name** command output, the fabric is partitioned into multiple CFS fabrics and the merge status may show that the merge has failed, is pending, or is waiting.

Merge Failure Troubleshooting

During a merge, the merge managers in the merging fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge. When a merge is successful, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. A merge failure indicates that the merged fabrics contain inconsistent data that could not be merged.

If a new switch is added to the fabric and the merge status for any application shows "In Progress" for a prolonged period of time, then there may be an active session for that application in some switch. Check the lock status for that application on all the switches using the **show cfs lock** CLI command. If there are any locks, then the merge will not proceed. Commit the changes or clear the session lock so that the merge can proceed.

Recovering from a Merge Failure with Fabric Manager

To recover from a merge failure using Fabric Manager, follow these steps:

Step 1 Select the **CFS** tab for the application that you are configuring and check the merge field to identify a switch that shows a merge failure. For example, choose **Fabricxx > All VSANS > DPVM** and select the **CFS** tab to determine if there is a merge failure for DPVM.

Step 2 Set the Config Action drop-down menu to **commit** and click **Apply Changes** to restore all peers in the fabric to the same configuration database.

Recovering from a Merge Failure with the CLI

To recover from a merge failure using the CLI, follow these steps:

Step 1 To identify a switch that shows a merge failure, issue the **show cfs merge status name** *application-name* command. Example command output follows:

```
Switch# show cfs merge status name ntp

Physical  Merge Status:Failure [ Mon Nov 22 06:49:52 2004 ]
 Local Fabric
-----------------------------------------------------------
 Switch WWN               IP Address
-----------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167   [Merge Master]
 20:00:00:0e:d7:00:3c:9e  10.76.100.52

 Remote Fabric
-----------------------------------------------------------
 Switch WWN               IP Address
-----------------------------------------------------------
 20:00:00:0d:ec:06:55:c0  10.76.100.205   [Merge Master]
```

Step 2 Enter configuration mode and issue the *application-name* **commit** command to restore all peers in the fabric to the same configuration database. Example command output follows:

```
Switch# config terminal
Switch(config)# ntp commit
Switch(config)#
```

Lock Failure Troubleshooting

In order to distribute a configuration in the fabric, a lock must first be acquired on all switches in the fabric. Once this is accomplished a commit can be issued which will distribute the data to all switches in the fabric before releasing the lock.

When a lock has been acquired by another application peer, you cannot commit new configuration changes. This is normal operation and you should postpone any changes to an application until the lock is released. Use the troubleshooting steps in this section only if you believe the lock has not been properly released.

A lock occurs when an administrator configures a change for a CFS-enabled application. If two administrators on the same switch attempt to configure the same application, only one administrator is given the lock. The other administrator is prevented from making changes to that application until the first administrator commits a change or discards any changes. Use the **show cfs lock name** CLI command to determine the name of the administrator who holds the lock for an application. You should check with that administrator before clearing the lock.

A CFS lock can also be held by another switch in your fabric. Use the **show cfs peers name** CLI command to determine all switches that participate in the CFS distribution for this application. That use the **show cfs lock name** CLI command on each switch to determine who owns the CFS lock for that applications. You should check with that administrator before clearing the lock.

Use the CFS **abort** option to release the lock without distributing the data to the fabric.

Resolving Lock Failure Issues Using Fabric Manager

To resolve a lock failure using Fabric Manager, follow these steps:

Step 1 Select the **CFS** tab for the application that you are configuring and view the **Master** check box to identify the master switch for that CFS application. For example, choose **Fabricxx > All VSANS > DPVM** and select the **CFS** tab.

Step 2 Set the Config Action drop-down menu on the master switch to **commit** or **abort** and click **Apply Changes** to restore all peers in the fabric to the same configuration database and free the CFS lock.

Resolving Lock Failure Issues Using the CLI

To resolve a lock failure using the CLI, follow these steps:

Step 1 Issue a **show cfs lock name** command to determine the lock holder. An example of the **show cfs lock name** command follows:

```
Switch# show cfs lock ntp
Application:ntp
Scope      :Physical
--------------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
--------------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3

Total number of entries = 1
```

Step 2 If the lock is being held by a remote peer, an *application-name* **commit** command or an *application-name* **abort** command must be executed at that switch. An example of the *application-name* **commit** command follows:

```
Switch# config terminal
Switch(config)# ntp commit
Switch(config)#
```

An example of the *application-name* **abort** command follows:

```
Switch# config terminal
Switch(config)# ntp abort
Switch(config)#
```

System State Inconsistent and Locks Being Held

An inconsistent system state occurs when locks are not held on all of the switches in the fabric, or when locks are held on all switches in the fabric, but a session does not exist with the lock holding switch. In either case, it may be necessary to use the **clear** option to release the locks.

Clearing Locks Using Fabric Manager

To clear a lock using Fabric Manager, follow these steps:

Step 1 Select the **CFS** tab for the application that you are configuring and view the **Master** check box to identify the master switch for that CFS application. For example, choose **Fabricxx > All VSANS > DPVM** and select the **CFS** tab.

Step 2 Set the Config Action drop-down menu on the master switch to **clear** and click **Apply Changes** to free the CFS lock.

Clearing Locks Using the CLI

When a lock is being held on a remote peer and issuing the *application-name* **commit** command or the *application-name* **abort** command does not clear the lock, issue the **clear** *application-name* **session** command to clear all locks in the fabric. After all locks are cleared, a new distribution must be started to restore all the switches in the fabric to the same state.

Example command output follows:

```
Switch# clear ntp session
Switch# config terminal
Switch(config)# ntp commit
Switch(config)#
```

Distribution Status Verification

After configuring an application and committing the changes, you may want to verify that CFS is distributing the configuration change throughout the fabric or VSAN.

Verifying Distribution Using Fabric Manager

In Fabric Manager, choose the **CFS** tab for the application that you are configuring and check the **Last Results** field to view the distribution status for your latest commit.

Verifying Distribution Using the CLI

In the CLI, use the **show cfs lock name** *application-name* command to determine if a distribution is in progress on the fabric. If the application does not show in the output, the distribution has completed. Example command output follows:

```
Switch# show cfs lock name ntp

Scope      :Physical
-----------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
-----------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3

Total number of entries = 1
```

CHAPTER 6

Troubleshooting Ports

This chapter describes how to identify and resolve problems that can occur with ports in the Cisco MDS 9000 Family of multilayer directors and fabric switches. It includes the following sections:

Overview

Before a switch can relay frames from one data link to another, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces (IPFC).

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, and B port. In addition to these modes, each interface may be configured in auto or Fx port modes. These modes determine the port type during interface initialization.

Each interface has an associated administrative configuration and operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (such as the operation speed).

For a complete description of port modes, administrative states, and operational states, refer to the *Cisco MDS 9000 Family Configuration Guide* and the *Cisco MDS 9000 Fabric Manager Configuration Guide.*

Best Practices

You can avoid potential problems by following best practices when you configure a port interface.

- Before you begin configuring a switch, make sure that the modules in the chassis are functioning as designed. Choose **Switches > Hardware** in Fabric Manager or use the **show module** CLI command to verify that a module is OK or active before continuing the configuration.
- Ensure that a Fibre Channel port is configured to the appropriate port mode for your configuration. The default port mode is auto on the 16-port 2-Gbps Fiber Channel switching modules and Fx on the 32-port 2-Gbps Fibre Channel switching modules.
- Configure devices attached to TL ports in zones.
- Observe the following guidelines when configuring a 32-port 2-Gbps Fibre Channel switching module or the Cisco MDS 9100 Series. When configuring these host-optimized ports, the following port mode guidelines apply:
  - You can configure only the first port in each 4-port group s an E port (for example, port 1 from ports 1-4, port 5 from ports 5-8, and so on). If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8, and so on) are not usable and remain shutdown.
  - If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
  - The auto mode is not allowed in a 32-port 2-Gbps Fibre Channel switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized Fibre Channel ports in the Cisco MDS 9120 switch and 32 host-optimized Fibre Channel ports in the Cisco MDS 9140 switch).
  - The default port mode is Fx (Fx negotiates to F or FL) for 32-port 2-Gbps Fibre Channel switching modules and the host-optimized Fibre Channel ports in the Cisco 9100 Series.
  - The 32-port 2-Gbps Fibre Channel switching module has not been qualified for FICON.

Initial Troubleshooting Checklist

Troubleshooting a SAN problem involves gathering information about the configuration and connectivity of individual devices and the entire SAN fabric. In the case of port interfaces, begin your troubleshooting activity as follows:

Checklist	Checkoff
Check the physical media to ensure there are no damaged parts.	❑
Verify that the SFP (small form-factor pluggable) devices in use are those authorized by Cisco and that they are not faulty.	❑
Verify that you have enabled the port by right-clicking the port in Device Manager and selecting **enable** or by using the **no shut** CLI command.	❑

Checklist (continued)	Checkoff
Right-click the port in Device Manager or use the **show interface CLI** command to verify the state of the interface. Refer to Table 6-1 for reasons why a port may be in a down operational state.	❑
Verify that you if you have one host-optimized port configured as an ISL, you have not connected to the other three ports in the port group.	❑

Table 6-1 Reason Codes for Nonoperational States

Reason Code	Description	Applicable Mode
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The Cisco SAN-OS software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state; then, administratively shut down and reenable the interface.	

Table 6-1 Reason Codes for Nonoperational States (continued)

Reason Code	Description	Applicable Mode
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to other side E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and RL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the isolation problem.

Limitations and Restrictions

- You must administratively enable a port with the **no shut** command. When the interface is enabled, the administrative state of the port is up. If you administratively disable an interface with the **shut** command, the administrative state of the port is down, and the physical link layer state change is ignored.
- For a port to be in an up operational state where it can transmit or receive traffic, the interface must be administratively up, the interface link layer state must be up, and the interface initialization must be complete.
- The interface cannot transmit or receive data when a port's operational state is down.
- The interface is operating in TE mode when a port's operational state is trunking.

Overview of the FC-MAC Driver and the Port Manager

This section describes the internal details of port related components in Cisco SAN-OS. Use this section to understand the underlying functions that may be causing port related problems.

The FC-MAC driver resides in the module component of the Cisco MDS 9000 Family SAN-OS software. It performs the following functions:

- Initialization of FC-MAC ASIC.
- Speed negotiation.
- Link/loop port initialization and credit recovery.
- Statistics collection.
- Error handling (mainly by acting on error interrupts).
- SFP detection and housekeeping.
- Statistics collection.
- Debug command support under the **show hardware internal fc-mac** command.

The FC-MAC driver does not handle FLOGI, RSCN, or configuration management.

This section includes the following topics:

Port Manager Overview

The Port Manager is management software running on the supervisor module. The Port Manager handles the following tasks:

- Port configuration management.
- Link events, including notifying the registered application on the supervisor module.

- E or TE port initialization.
- SFP validation.

The FC-MAC detects the port is in one of the following states:

- Disable—The port is administratively disabled.
- Enable—The port is administratively enabled. In this state, the port may be in speed initialization, loop-initialization, link (point-to-point connection) initialization, or the link-up state.
- HW Failure—The port has been declared bad due to a hardware failure.
- Pause—An intermediate state after the link is down and subsequent enabling of the port to start the port initialization.

You can check the state of the port using the command:

show hardware internal fc-mac port *slot/port* **port-info**

The FLOGI server is a separate application that handles the FLOGI processing for Nx ports.

Troubleshooting Port States with the Device Manager

Device Manager offers three ways to monitor ports:

- Device View
- Summary View
- Port Selection

Device View

Basic port monitoring using Device Manager begins with the visual display in the Device View (Figure 6-1). Port display descriptions include:

- Green box—A successful fabric login has occurred; the connection is active.
- Red X—A small form-factor pluggable transceiver (SFP) is present but there is no connection. This could indicate a disconnected or faulty cable, or no active device connection.
- Red box—An FSP is present but fabric login (FLOGI) has failed. Typically a mismatch in port or fabric parameters with the neighboring device. For example, a port parameter mismatch would occur if a node device were connected to a port configured as an E port. An example of a fabric parameter mismatch would be differing timeout values.
- Yellow box—In Device Manager, a port was selected.
- Gray box—The port is administratively disabled.
- Black box—FSP is not present.

Figure 6-1 Device Manager: Device View

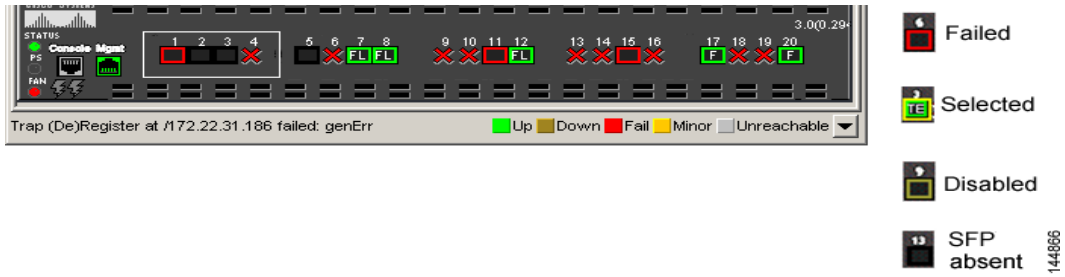

Device Manager: Summary View

In Device Manager, selecting the Summary View (Figure 6-2) expands on the information available for port monitoring. The display includes:

- VSAN assignment
- For N ports, the port world-wide name (pWWN) and Fibre Channel ID (FC ID) of the connected device
- For ISLs, the IP address of the connected switch
- Speed
- Frames transmitted and received
- Percent utilization for the CPU, dynamic memory, and Flash memory

Figure 6-2 Device Manager: Summary View

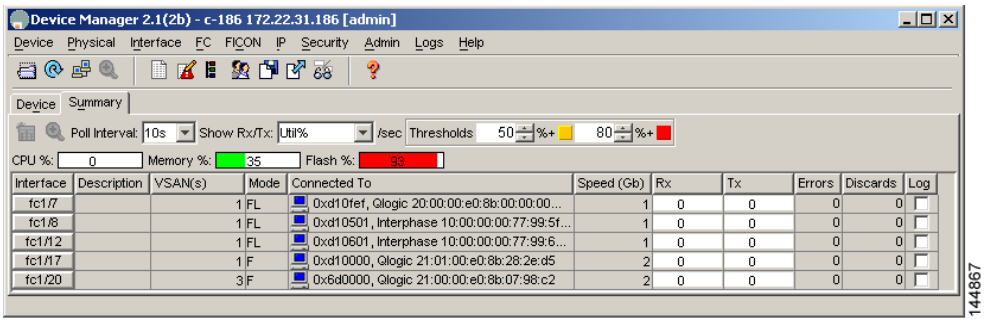

Device Manager: Port Selection

To drill down for additional port information, using either the Device view or Summary view, select and double-click any port. The initial display (Figure 6-3) shows administrative settings for Mode, Speed, and Status, plus current operational status, failure cause, and date of the last configuration change.

Additional tabs include:

- Rx BB Credit–Configure and view buffer-to-buffer credits (BB credits).
- Other–View PortChannel ID, WWN, Maximum Transmission Unit (MTU), configure maximum receive buffer size.
- FLOGI–View FC ID, pWWN, nWWN, BB credits and class of service for N port connections.
- ELP–View pWWN, nWWN, BB credits and supported classes of service for ISLs.
- Trunk Config–View and configure trunk mode and allowed VSANs.
- Trunk Failure–Failure cause for ISLs.
- Physical–Configure beaconing; view SFP information.
- Capability–View current port capability for hold-down timers, BB credits, maximum receive buffer size.

Figure 6-3 Device Manager: Port Selection

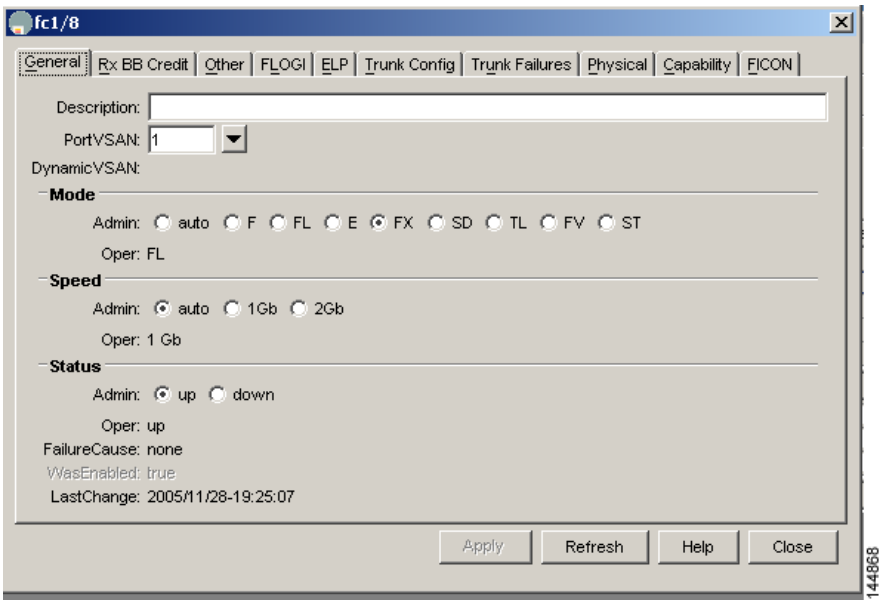

Troubleshooting Port States from the CLI

To display complete information for an interface, use the show **interface** CLI command. In addition to the state of the port, this command displays:

- Port WWN
- Speed
- Trunk VSAN status
- Transmit and receive buffer-to-buffer credits configured and remaining

- Maximum receive buffer size
- Number of frames sent and received
- Transmission errors, including discards, errors, CRCs, and invalid frames

Example 6-1 displays the **show interface** CLI command output.

Example 6-1 show interface Command Output

```
    switch# show interface fc1/3
    fc1/3 is trunking
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:03:00:0b:fd:8c:f8:80
    Peer port WWN is 20:10:00:0b:fd:2c:8c:00
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 161
    Speed is 2 Gbps
    Transmit B2B Credit is 255
Receive B2B Credit is 255
Receive data field Size is 2112
```

Using Port Debug Commands

Use the **show hardware internal debug-info interface fc** CLI command to debug ports.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

Examples of when to use these commands include:

- An Fibre Channel port fails to move to the up state after such events as link failures, admin-up operations, or new connections.
- Unexpected link flaps.
- The port moves to “error disabled” state.

Maintain a set of information for the module before these problems occur (if possible) and then gather another set of information after these problems occur.

Useful Commands at the FC-MAC Level

Troubleshooting a port problem involves analysis of the debug facilities provided by the FC-MAC driver, or the FC-MAC2 driver in the case of the MDS 9120, MDS 9140, MDS 9216i, and the MPS-14/2 module. Table 6-2 lists several CLI debugging commands at the FC-MAC level.

Note Use the **fcmac2** keyword for the MDS 9120, MDS 9140, MDS 9216i, and the MPS-14/2 module.

Table 6-2 Useful FC-MAC Port Commands

CLI Command	Description
show hardware internal fc-mac port *slot/port* **link-status**	Performs a series of checks to isolate the problem.
show hardware internal fc-mac2 port *slot/port* **link-status**	
show hardware internal fc-mac port *slot/port* **port-info**	Provides the current state and configuration of the port.
show hardware internal fc-mac2 port *slot/port* **port-info**	
show hardware internal fc-mac port *slot/port* **statistics**	Gives all non-zero statistics for the port.
show hardware internal fc-mac2 port *slot/port* **statistics**	
show hardware internal fc-mac port *slot/port* **gbic-info**	Displays the current state of the SFP.
show hardware internal fc-mac2 port *slot/port* **gbic-info**	
show hardware internal error	Collects interrupt statistics, error statistics, and exception log information for the entire module.
show hardware internal debug-info interface *fc-interface*	Represents an aggregation of a number of debug commands from all ASICs. The information includes interrupt-statistics, error-statistics, exception-log, link-events, and all debug information that is provided by the FC-MAC driver.

Note To issue CLI commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

Common Problems with Port Interfaces

The following issues are commonly seen with port interfaces:

Port Remains in a Link Failure or Not Connected State

If a link does not come up, then the switch was unable to achieve bit or word synchronization with the node device. This situation may occur if nothing is connected to the interface, as in the case of a broken fibre, or if there is no bit synchronization between the switch interface and the directly connected Nx port. This problem may be the result of one or more of the possible causes listed in Table 6-3.

Symptom Port remains in a link-failure state.

Table 6-3 Port Remains in a Link-Failure State

Symptom	Possible Cause	Solution
Port remains in a link-failure state.	Port connection is bad.	Use the **show port internal info** CLI command to verify the port status is in link-failure. Use the **show hardware internal fc-mac port** *slot /port* **gbic-info** CLI command to determine if there is a signal present. Verify the type of media in use. Is it copper or optical, single-mode (SM) or multimode (MM)? Verify that the media is not broken or damaged. Is the LED on the switch green? Is the active LED on the host bus adapter (HBA) for the connected device on? Right-click on the port in Device Manager and select **disable** and then **enable**, or use the **shut CLI** command followed by the **no shut** command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.
	There is no signal because of a transit fault in the SFP or the SFP may be faulty.	When this occurs, the port stays in a transit port state and you see no signal. There is no synchronization at the MAC level. The problem may be related to the port speed setting or autonegotiation. See the "Troubleshooting Port Problems" section on page 6-12. 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.
	Link is stuck in initialization state or the link is in a point-to-point state.	Choose **Logs > Switch Resident > Syslog** on Device Manager or use the **show logging** CLI command to check for a `Link Failure, Not Connected` system message. Right-click on the port in Device Manager and select **disable** and then **enable**, or use the **shut CLI** command followed by the **no shut** command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem.

Troubleshooting Port Problems

Start the debugging with the command **show hardware internal fc-mac port** *slot/port* **link-status**. See the "Useful Commands at the FC-MAC Level" section on page 6-9 to understand how to use the FC-MAC information. When this command executes, it performs the following checks in the order shown here and displays the appropriate information:

1. Checks whether the port was declared a failure because of an exception. For additional information, use the **show process exceptionlog** CLI command.
2. Checks whether the port is administratively enabled.
3. Checks whether the physical link state is up. If the state is up, then it does the following:
   - Checks for possible completion of the FLOGI process.

     Note FLOGI is transparent to the MAC driver and is based on some expected configuration. The MAC driver assumes that the FLOGI process is completed.

   - Checks for error counters.
4. Checks whether the port is in the offline state. The port goes to the offline state if the FLOGI or ELP (in case of auto mode) on the port does not succeed.
5. Checks for pause state. A pause state is in an intermediate state (as maintained by the FC-MAC driver) after the link goes down and before the port is enabled by the Port Manager.

Note The link reinitializes after a link down event is initiated only if enable is issued by the Port Manager.

6. Checks for the presence of SFP/GBIC. If present, FC-MAC checks for loss of signal. The loss of signal state indicates either the physical connectivity between two end ports is bad or there is a transmit fault in the SFP. Use the **show hardware internal fc-mac port** *slot/port* **gbic-info** command to check for the transmit fault.
7. Checks for the speed and sync state of the port. If the port is in the speed initialization state, then:
   - `Auto speed is in progress` is displayed if the port is in automode.
   - `Waiting for stable sync` is displayed if the port is configured for a fixed speed.
   - `Sync not acquired` is displayed if the MAC state indicates a loss of synchronization. In auto mode, this state is not necessarily an error. In any case, check the speed capabilities and configuration at both ends.

Port Remains in Initializing State

Symptom Port remains in the initializing state.

A port goes into the initialization state after a successful completion of link level initialization. For Fx and FL types of ports, the next step is to complete the FLOGI process. The port remains in the initialization state until the FLOGI (fabric login) process completes.

For E or TE port types, the next step is to complete the ELP process. If the ELP fails the port is moved to the offline state after a timeout and the entire process repeats until the port comes online.

Table 6-4 lists possible causes for FLOGI to fail for a given port and possible solutions.

Table 6-4 Port Remains in the Initializing State

Symptom	Possible Cause	Solution
Port remains in the initializing state.	The port is up because the link partner has put itself in a bypass mode.	Use the **show hardware internal fc-mac port** *slot/port* **statistics** command to check whether the Class-3 input counter is increasing after the successful completion of link initialization.
	The FLOGI packet was dropped somewhere in the data path, starting from FC-MAC to the FLOGI server.	Use the **show hardware internal fc-mac port** *slot/port* **statistics** command to check for Class-3 packet counters. Analyze the output of the **show hardware internal error** command for a possible drop of FLOGI packets somewhere in the path. See the "NoteWe recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem." section on page 6-13. Right-click on the port in Device Manager and select **disable** and then **enable**, or use the **shut CLI** command followed by the **no shut** command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.
	A software bug resulted in an error while handling the FLOGI packet.	

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem.

Troubleshooting Port Registration Issues Using the CLI

To troubleshoot Nx port registration in the CLI, follow these steps:

Step 1 Use the **show interface fc** *slot/port* CLI command and verify that the fibre channel interface connected to the device in question is up and free of any errors. (See Example 6-2.)

Example 6-2 show interface Command Output

```
switch# show interface 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:

show flogi database vsan *vsan-id*

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 7. If the device does not appear in the output, go to the next step.

Step 3 Use the **shutdown** CLI command in interface configuration mode to shut down the Fibre Channel interface connected to the device in question.

```
switch# config terminal
switch(config)# interface fcx/x
switch(config-if)# shutdown
```

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem.

Step 4 Use the **no shutdown** CLI command on the Fibre Channel interface.

```
switch(config-if)# 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.

Use the **show flogi internal event-history interface** CLI command to view the events that occurred on the interface after you enabled it again. The comments that follow each section of output explain the meaning of the output.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
switch# show flogi internal event-history interface fc3/14
>>>>FSM: <[99]21:00:00:e0:8b:07:a4:36> has 9 logged transitions<<<<<
/* This is the [VSAN] followed by the pwwn 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 */

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_RESPOSE]
    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 */
```

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 5 Use the **shutdown** CLI command in interface mode to shut down the Fibre Channel interface Then use the **no shutdown** CLI command after turning on the debug described in Step 6 and Step 7.

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem.

Step 6 Use the **debug fcns events register vsan** CLI command to watch the FLOGI process take place.

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

This command enables debug mode for name server registration. It generates messages on the switch console related to FCNS events. The system output may look something like this:

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc3/14
switch(config-if)# no shutdown   /* 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 */

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 similar to these will be listed if more name server objects are registered.

Step 7 If you are managing the switch over a Telnet connection, enable terminal monitoring by entering the **terminal monitor** CLI command in exec mode.

The system output looks like this:

```
switch# show fcns database detail vsan 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.

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

Unexpected Link Flapping Occurs

Symptom Unexpected link flapping occurs.

When a port is flapping, it cycles through the following states, in this order, and then starts over again:

1. Initializing - The link is initializing.
2. Offline - The port is offline.
3. Link failure or not connected - The physical layer is not operational and there is no active device connection.

When troubleshooting unexpected link flapping, it is important to know the following information:

- Who initiated the link flap.
- The actual link down reason.

Be sure to check the HBA, because a faulty HBA can manifest symptoms on the attached switch port. For example, if an Nx port is self-diagnosed as faulty by the HBA driver or firmware, the driver can place the port in optical bypass mode. This results in the receive and transmit paths being internally connected through the port. 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 to try to initialize as an xE port, even if an end device is physically connected to that interface. In this case, a port reason code of isolation because of ELP failure can be displayed even if an ISL is not present.

Table 6-5 lists possible causes and solutions for link flapping.

Table 6-5 Unexpected Link Flapping Occurs

Symptom	Possible Cause	Solution
Unexpected link flapping occurs.	The bit rate exceeds the threshold and puts the port into an error disabled state.	Right-click the port in Device Manager and select **disable** and then **enable**, or use the **shut CLI** command followed by the **no shut** command to return the port to the normal state.
	The switch cannot complete the link reset. The link reset protocol failure results in a link flap that may be the result of: • The input buffer did not become empty within the link reset timeout period. • The link partner did not respond to a link reset initiated by the switch.	The switch initiates the link reset when all credits are lost for more than four seconds or when there is a temporary signal or sync loss condition that lasts for less than 100msec. See the "Troubleshooting Port Problems" section on page 6-12 to verify this condition. Right-click the port in Device Manager and select **disable** and then **enable**, or use the **shut CLI** command followed by the **no shut** command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.
	There is a credit loss condition on an FL port.	When credit loss or a transmit stuck condition is detected in the FL port, the FC-MAC drive flaps the link as a recovery process. See the "Troubleshooting Port Problems" section on page 6-12.
	Some problem in the switch triggers the link flap action by the end device. Some of the causes are: • Packet drop in the switch, because of either a hardware failure or an intermittent hardware error such as X-bar sync loss. • Packet drop resulting from a software error. • A control frame is erroneously sent to the device.	Determine link flap reason as indicated by the MAC driver . Use the debug facilities on the end device to troubleshoot the problem. An external device may choose to reinitialize the link upon encountering the error. In such cases, the exact method of reinitializing the link varies by device. See the "Troubleshooting Port Problems" section on page 6-12 for more information on externally triggered link flaps.

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the problem.

Link Initialization Flow

Fibre Channel primitive sequences are used to establish and maintain a link and they continue to be transmitted until a response has been received. Four primitive sequences are used in the link initialization process:

- Not operational sequence (NOS)
- Offline sequence (OLS)
- Link reset sequence (LRS)
- Link reset response sequence (LRR)

Figure 6-4 uses the ordered sets of 8b/10 encoding in the primary operational states. They include:

- AC = Active state
- LR = Link recovery state
- LF = Link failure state
- OF = Offline state

Figure 6-4 Link Initialization Flow

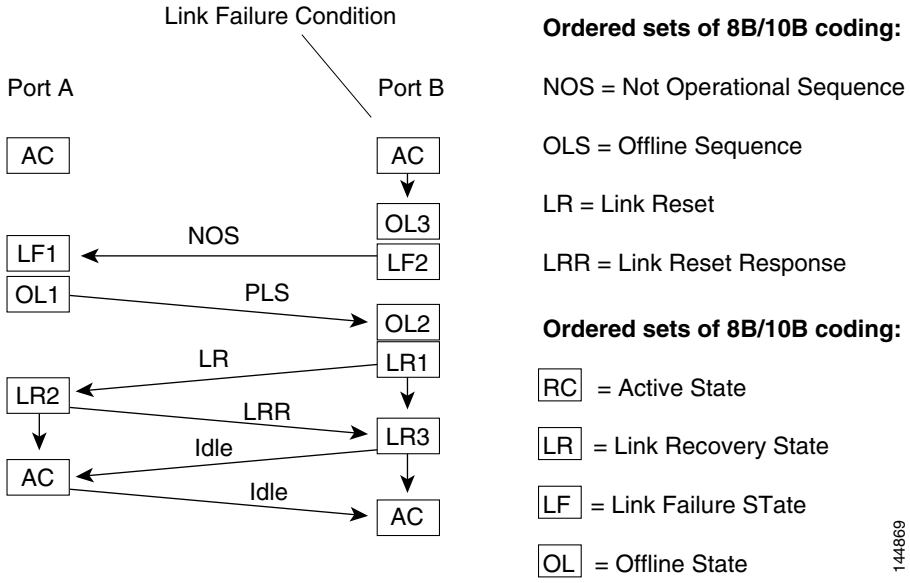

Figure 6-4 shows the link initialization flow. It displays the ordered sets transmitted between the ports and the primary operational states of the port during the process. They include:

1. Active state.
2. Link recovery state (LR):
   a. LR transmit substate (LR1)
   b. LR receive substate (LR2)
   c. LRR receive substate (LR3)

3. Offline state (OLS):
   a. OLS transmit substate (OL1)
   b. OLS receive substate (OL2)
   c. Wait for OLS substate (OL3)
4. Link failure state:
   a. NOS receive substate (LF1)
   b. NOS transmit substate (LF2)

The Cisco MDS 9000 Family switch maintains port counters for link initialization ordered sets, including OLS, LRR, and NOS for fabric connections, as well as primitives for arbitrated loop connections on FL ports and TL ports. Understanding the link initialization flow and viewing the port counters using **show interface** can be useful when you troubleshoot port initialization problems. Table 6-6 displays the reasons for a link flap.

Table 6-6 Link Flap Reasons Initiated by a Device Connected to the Switch Port

Reason	Description
Sync Loss	A synchronization loss condition persisted for more than 100 milliseconds. Look at the Invalid Transmission Word Count to check whether the physical link is really bad and if that caused the loss of synchronization. Sometimes this is not necessarily a problem with the physical link, but with the way some devices initialize the link. Use **attach module** to connect to the module and then use the **show hardware internal debug-info interface** CLI command. See Table 6-2.
Loss of signal	A signal loss condition persisted for more than 100 milliseconds. Look at the Invalid Transmission Word Count to check whether the physical link is really bad and if that caused the loss of synchronization. Sometimes this is not necessarily a problem with the physical link, but with the way some devices initialize the link. If the link does not come up after a flap, then probably the other end is in a shutdown state or the cable is broken. You can check for the broken or disconnected optical link by using the **show hardware internal fc-mac port** *slot/port* **gbic-info** CLI command.
NOS received	A NOS received condition is detected. If the other end is an MDS port, then the NOS is transmitted by the other end in one of the following conditions: • A signal loss or sync loss condition is detected. • The port is administratively shut down. • The port is operationally down.
OLS received	An OLS received condition is detected.
LR received B2B	Link reset (LR) failed because of the receive queue (in the queue engine) not being empty.
Cr loss	Too many credit loss events occurred.
Rx queue overflow	The receive queue overflowed in the queue engine occurred. This can happen under the following conditions: • Improper credit configuration at one or both ends of the link. • A bad link can sometimes result in extra R_RDYs. Check for invalid transmission words at both ends.

Table 6-6 Link Flap Reasons Initiated by a Device Connected to the Switch Port (continued)

Reason	Description
LIP F* received	An loop initialization procedure (LIP) was received.
LC port shutdown	The port shutdown was invoked. Use the **show process exception** CLI command to check for any other errors.
LIP received B2B	An LIP was received while the Rx queue was not empty.
OPNy tmo B2B	An open circuit on a loop (OPNy) timeout occurred while the Rx queue was not empty.
OPNy Ret B2B	An OPNy was returned while the Rx queue was not empty.
Cr Loss B2B	Credit loss occurred while the Rx queue was not empty.

Viewing Port Counters

You can use the **show interface counters** command to view port counters. Typically, you only observe counters while actively troubleshooting, in which case you should first clear the counters to create a baseline. The values, even if they are high for certain counters, can be meaningless for a port that has been active for an extended period. Clearing the counters provides a better idea of the link behavior as you begin to troubleshoot.

Use one of the following commands in EXEC mode to clear all port counters or counters for specified interfaces:

- **clear counters interface all**
- **clear counters interface** <*range*>

The counters can identify synchronization problems by displaying a significant disparity between received and transmitted frames. For example, in the case of a broken fiber, if only the Tx path from the F port to the N port is broken, then the switch interface will still have an operational Rx path and will still obtain bit synchronization from the bit stream received from the N port. The switch port will also be able to recognize an incoming NOS from the N port and reply with an OLS. However, 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 and NOS increment (as there is bit synchronization but no word synchronization). In such a state, you can check that the Tx path from the switch to the Rx input on the 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.

The output in Example 6-3 also displays evidence of corrupt data on the wire if there are a high number of CRCs and errors. Discards may or may not indicate a problem. For example, a frame can be discarded because of an ACL violation.

Example 6-3 show interface Command

```
mds# show interface fc4/2
fc4/2 is up
  . . .
    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
      322944 frames input, 19378384 bytes
      0 discards, 0 errors   <..................................... Errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      20439797 frames output, 41780390808 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      2 output OLS, 2 LRR, 0 NOS, 0 loop inits <.................Link Initialization
      12 receive B2B credit remaining
      1 transmit B2B credit remaining
```

Port Bounces Between Initializing and Offline States

Symptom Port bounces between the initializing and offline states.

An ELP failure may result in a port bouncing between the initializing and offline states. Table 6-7 lists possible causes and solutions to this problem.

Table 6-7 Port Bounces Between the Initializing and Offline States

Symptom	Possible Cause	Solution
Port bounces between the initializing and offline states.	An ELP packet was dropped in one of the two switches.	Use the **show hardware internal fc-mac port** *slot/port* **statistics** CLI command and the **show hardware internal error** command. Analyze the output of the two commands for possible packet drops. See the "Troubleshooting ELP Issues Using the CLI" section on page 6-24. See also the "xE Port Is Isolated in a VSAN" section on page 7-7.
	There is a software bug or incompatibility in handling the ELP process.	Analyze the event history provided by the Port Manager after using the **show port internal event-history** CLI command. See the "Troubleshooting ELP Issues Using the CLI" section on page 6-24.

Troubleshooting ELP Issues Using the CLI

To troubleshoot ELP issues using the CLI, follow these steps:

Step 1 **Use the show interface CLI command to verify E port isolation**:

```
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
    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 indicates 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.

Step 2 Verify that the following parameters match on each switch in the VSAN using the **show fctimer** CLI command:

- ED_TOV timer
- RA_TOV timer
- FS_TOV timer

Note Because fabric parameters are configured on a per VSAN basis, they are required to be the same for all switches within a VSAN.

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

This sample output shows the default settings for these timeout values.

Step 3 Optionally, use the **fctimer** CLI command in config mode to globally set these timeout values across all VSANs or use the **fctimer D_S_TOV** <*timeout*> **vsan** <*vsan-id*> CLI command for example, to set the D_S_TOV timeout for a particular VSAN to override the global values.

Step 4 Use the **show port internal info interface fc** CLI command to verify that Rx buffer size matches on both ends of the ISL.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
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 Sourcable Pkt Size: 2112
    Hw Capabilities: 0xb
    Connector Type: 0x0
  SFP 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)
```

E Port Bounces Remains Isolated After a Zone Merge

Symptom E port remains isolated after a zone merge.

An E port may be isolated because of a zone merge failure. Table 6-8 lists possible causes and solutions to this problem.

Table 6-8 E Port Remains Isolated after a Zone Merge

Symptom	Possible Cause	Solution
E port remains isolated after a zone merge.	The active zone sets on the two switches differ from each other in terms of zone membership (provided there are zones at either side with identical names).	See the "Troubleshooting E port Isolation using Fabric Manager" section on page 6-26 or the "Troubleshooting E port Isolation using Fabric Manager" section on page 6-26.
	The active zone set on both switches contains a zone with the same name but with different zone members.	

Troubleshooting E port Isolation using Fabric Manager

To troubleshoot E port isolation due to zoning using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > FC Physical** to verify that the E port did not come up because of a zone merge failure.

Note Zoning information exists 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.

Step 2 Select **Zone > Edit Local Full Zone Database** to verify the zoning configuration.

Step 3 Use one of the following two approaches to resolve a zone merge failure:

- Choose **File > Restore** from the Edit Local Full Zone Database dialog box to overwrite the zoning configuration of one switch with the other switch's configuration.

  The **Restore** option overwrites the local switch's active zone set with that of the remote switch.
- If the zoning databases between the two switches are overwritten, you cannot use the **Restore** option. To work around this, you can manually change the content of the zone database on either of the switches using the Edit Local Full Zone Database, and then choose **Switches > Interfaces > FC Physical** and select **down** and then **up** on the Admin Status drop-down menu for the isolated port.

Step 4 If the isolation is specific to one VSAN and not on an E port, the correct way to issue the cycle up or down is to remove the VSAN from the list of allowed VSANs on that trunk port, and reinsert it.

a. Choose **Switches > Interfaces > FC Physical** and select the **Trunk Config** tab.

b. Remove the VSAN from the Allowed VSAN list and click **Apply Changes.**

c. Add the VSAN back to Allowed VSAN list and click **Apply Changes.**

Note We recommend that you do not disable and then enable a T or TE port. 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 zone sets in two switches can be checked before actually merging the two zone sets. Refer to the *Cisco MDS 9000 Fabric Manager Configuration Guide* for more information.

Troubleshooting E port Isolation Using the CLI

To troubleshoot E port isolation due to zoning using the CLI, follow these steps:

Step 1 Use the **show interface** command output to verify that the E port did not come up because of a zone merge failure.

Note Zoning information exists 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
```

Step 2 Verify the zoning information using the following commands:

- **show zone vsan** *vsan-id*
- s**how zoneset vsan** *vsan-id*

Step 3 Use one of the following two approaches to resolve 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:
  - **zone copy interface fc** *slot/port* **import vsan** *vsan-id*
  - **zone copy interface fc s***lot/port* **export vsan** *vsan-id*

  The **import** option of the command of overwrites the local switch's active zoneset with that of the remote switch. The **export** option overwrites the remote switch's active zoneset with the local switch's active zone set.
- If the zoning databases between the two switches are overwritten, you cannot use the **import** option. To work around this, you can manually change the content of the zone database on either of the switches, and then issue a **shutdown/no shutdown** command sequence on the isolated port.

Step 4 If the isolation is specific to one VSAN and not on an E port, the correct way to issue the cycle u p or down is to remove the VSAN from the list of allowed VSANs on that trunk port, and reinsert it.

Note We recommend that you do not disable and then enable a T or TE port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the isolation problem.

Port Cycles Through Up and Down States

Symptom Port cycles through the up and down states.

This problem may be attributable to an error experienced by the connected device. Table 6-9 lists the possible causes and solutions for this problem.

Table 6-9 Port Cycles Through the Up and Down States

Symptom	Possible Causes	Solutions
Port cycles through the up and down states.	One or more packets were dropped in the switch.	Analyze the debug log provided by the Nx port. Select **Tools > Traceroute** using Fabric Manager or use the **fctrace** CLI command to analyze the the link. Look for FLOGI messages in the logs for this port. See the "Troubleshooting Port Registration Issues Using the CLI" section on page 6-14
	There is a problem in FLOGI processing.	
	The device received unexpected packets.	
	There was a higher layer software error.	

Port Is in ErrDisabled State

The ErrDisabled state indicates that the switch detected a problem with the port and disabled the port. This state could be caused by a flapping port or a high amount of bad frames (CRC errors), potentially indicating something wrong with the media.

Symptom Port is in ErrDisabled state.

An E port may be isolated because of a zone merge failure. Table 6-10 lists possible causes and solutions to this problem.

Table 6-10 Port is in ErrDisabled State

Symptom	Possible Cause	Solution
Port is in ErrDisabled state.	Flapping port.	See the "Verifying the ErrDisable State Using the CLI" section on page 6-29. Verify the SFP, cable, and connections.
	Switch detected a high amount of bad frames (CRC errors), potentially indicating something wrong with the media.	

Verifying the ErrDisable State Using the CLI

To resolve the ErrDisable state using the CLI, follow these steps:

Step 1 Use the **show interface** command to verify that the switch detected a problem and disabled the port. Check cables, SFPs, and optics.

```
mds# show interface fc1/14
fc1/14 is down (errDisabled)
```

Step 2 Use the **show port internal event-history interface** command to view information about the internal state transitions of the port. In this example, port fc1/7 entered the ErrDisabled state because of a capability mismatch, or "CAP MISMATCH." You might not know how to interpret this event, but you can look for more information with other commands.

```
mds# show port internal event-history interface fc1/7
>>>>FSM: <fc1/7> has 86 logged transitions<<<<<
1) FSM:<fc1/7> Transition at 647054 usecs after Tue Jan  1 22:44..
    Previous state: [PI_FSM_ST_IF_NOT_INIT]
    Triggered event: [PI_FSM_EV_MODULE_INIT_DONE]
    Next state: [PI_FSM_ST_IF_INIT_EVAL]
2) FSM:<fc1/7> Transition at 647114 usecs after Tue Jan  1 22:43..
    Previous state: [PI_FSM_ST_IF_INIT_EVAL]
    Triggered event: [PI_FSM_EV_IE_ERR_DISABLED_CAP_MISMATCH]
    Next state: [PI_FSM_ST_IF_DOWN_STATE]
```

Step 3 Use the **show logging logfile** command to display the switch log file and view a list of port state changes. In this example, an error was recorded when someone attempted to add port fc1/7 to PortChannel 3. The port was not configured identically to PortChannel 3, so the attempt failed.

```
mds# show logging logfile
 . . .
Jan  4 06:54:04 switch %PORT_CHANNEL-5-CREATED: port-channel 17 created
Jan  4 06:54:24 switch %PORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
17 is down (No operational members)
Jan  4 06:54:40 switch %PORT_CHANNEL-5-PORT_ADDED: fc1/8 added to port-channel 7
Jan  4 06:54:56 switch %PORT-5-IF_DOWN_ADMIN_DOWN: Interface fc1/7 is down
(Admnistratively down)
Jan  4 06:54:59 switch %PORT_CHANNEL-3-COMPAT_CHECK_FAILURE: speed is not compatible
Jan  4 06:55:56 switch %PORT_CHANNEL-5-PORT_ADDED: fc1/7 added to port-channel 7
```

Troubleshooting Fx Port Failure

Fx port problems can be caused by a variety of configuration issues. While most issues can be solved by simply ensuring that the ports are configured properly, some issues require the use of more in-depth troubleshooting techniques.

Overview of Symptoms

An F port may be connected to a single N port, which is the mode used by peripheral devices (hosts or storage). In all the possible cases an administrator can encounter in troubleshooting an Fx port, two different scenarios can be recognized:

- The port does not come up (check the interface configuration, cabling and the port connected to the switch).
- The port comes up, but the host cannot communicate with the storage subsystem (check the VSAN and zone configurations).

Typical end-user questions that lead to Fx port troubleshooting include:

- Why is no storage visible on my newly installed server?
- Why is previously assigned storage not visible to my server after reboot?

Typical administrator questions to investigate:

- Why does the server fail to complete FLOGI to the switch?
- Why does the storage device fail to complete FLOGI to the switch?

Figure 6-5 illustrates one possible methodology for troubleshooting Fx ports.

Figure 6-5 Troubleshooting Methodology

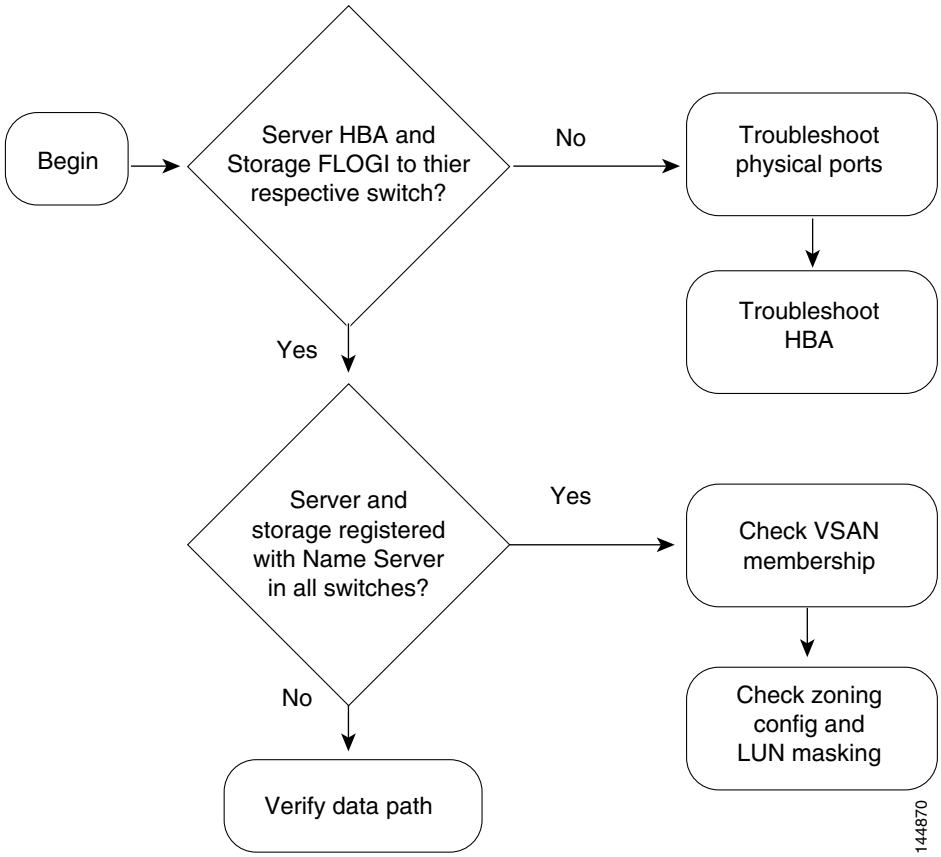

CHAPTER 7

Troubleshooting VSANs, Domains, and FSPF

This chapter describes how to identify and resolve problems that might occur when implementing VSANs, domains, and FSPF. This chapter includes the following sections:

Best Practices for VSAN Implementation

Virtual SANs (VSANs) provide a method of isolating devices that are physically connected to the same storage network, but are logically considered to be part of different SAN fabrics that do not need to be aware of one another. VSANs provide a way to:

- Isolate devices physically connected to the same fabric.
- Reduce the size of a Fibre Channel distributed database.
- Enable more scalable and secure fabrics.

This section provides the best practices for implementing VSANs.

- Avoid using VSAN 1 (the default VSAN) for production network traffic. Create at least one VSAN to carry your network traffic.
- Isolate devices in VSANs whenever practical.
- Leave fabric timers and FSPF timers at their default settings.

  Avoid modifying fabric timers and FSPF timers unless changes are required because of interoperability with an existing fabric, or long-haul links are being deployed.

- Use Inter-VSAN routing (IVR) only when necessary to selectively connect devices across VSANs.
  - If IVR is used without NAT, ensure that domain IDs are are statically configured and unique across all VSANs.
- Place FCIP gateways in their own native VSAN.

  Placing FCIP gateways in their own VSAN isolates disturbances when problems in the IP cloud (such as flapping links) occur.
- Use VSAN-based roles to control and limit management access to your switches.

Best Practices for Domain ID Assignment

This section provides best practices for domain ID assignments.

- Use static domains in most environments. To use static domains, choose **Fabric***xx* **> All VSANs > Domain Manager** and select **static** from the Config Type drop-down menu in Fabric Manager or use the **fcdomain domain** *n* **static vsan** *x* CLI command. You must then issue a disruptive restart so that the configured domain ID matches the running domain ID. Select the **Configuration** tab and select **disruptive** from the Restart drop-down menu in Fabric Manager and click **Apply Changes**. In the CLI, use the **fcdomain restart disruptive** CLI command.

Note You cannot issue a disruptive restart for VSANs that are in any of the interop modes. Use a nondisruptive restart as needed.

- To disable the Domain manager, choose **Fabric***xx* **> All VSANs > Domain Manager** and uncheck the **Enable** check box in Fabric Manager or use the **no fcdomain vsan** *x* CLI command.
  - Disable the Domain Manager to disable the principal switch selection process. This is possible if all domains are statically assigned. Disabling principal switch selection can reduce disruption when switches are rebooted or added to the fabric. This must be done on each switch that should not participate in principal switch selection. A disruptive restart of the fabric is required to apply this change.
- Keep domain ID allowed lists the same on all switches in a fabric for consistency. If the principal switch changes, the allowed domain lists will remain the same.
- Assign domain IDs between decimal 97 and 127 if the domain may be used for standards-based interop mode.
- Do not perform frequent changes to the Domain Manager on production fabrics. Experienced administrators familiar with switch operations should be responsible for Domain Manager changes. Plan your domain configuration carefully so that you avoid the need to make disruptive changes at a later time.

- Save Domain Manager changes. When you change the configuration, be sure to save the running configuration by choosing **Switches > Copy Configuration** in Fabric Manager or using the **copy running-config startup-config** CLI command. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.
- Enable reconfigure fabric (RCF) rejection on every ISL port if high availability is mandatory. Choose **Switches > Interfaces > FC Physical** in Fabric Manager and select the **Domain Manager** tab in the Information pane and then check the **RcfReject** check box on all ISL ports to enable rcf-rejects. Or use the **interface** CLI command on a TE or E port and then use the **fcdomain rcf-reject vsan** CLI command in interface configuration mode to enable the rcf-reject option. RCF reject prevents other switches from sending an RCF and potentially causing a disruption in your production traffic.

Best Practices for FSPF

This section provides best practices for implementing FSPF.

- Use the default FSPF link cost, which can be configured on a per-VSAN basis for the same physical link, provides preferred and alternate paths. If you must alter the FSPF link cost, use caution to avoid the potential for asymmetric Fibre Channel routing.
- Use the default FSPF load-balancing configuration unless you are required to load balance based on your unique fabric, for example, if you have FICON VSANs.
- Use the default FSPF timer configuration. If FSPF timers are misconfigured, then the switches will not reach the “two-way” state and FSPF will not operate properly.

Initial Troubleshooting Checklist

Most VSAN problems can be avoided by following the best practices for VSAN implementation. In addition to Fabric Manager and the CLI, 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 configuration consistency check tool is also provided by Fabric Manager. Refer to the *Cisco MDS 9000 Fabric Manager Configuration Guide* for more information about this tool.

Troubleshooting a SAN problem involves gathering information about the configuration and connectivity of individual devices and the entire SAN fabric. In the case of VSANs, begin your troubleshooting activity as follows:

Checklist	Checkoff
Verify the FSPF parameters for switches in the VSAN.	❑
Verify the domain parameters for switches in the VSAN.	❑
Verify the physical connectivity for any problem ports or VSANs.	❑
Verify that you have both devices in the name server.	❑
Verify that you have both end devices in the same VSAN.	❑

Checklist (continued)	Checkoff
Verify that you have both end devices in the same zone.	❑
Verify that the zone is part of the active zone set.	❑

Common Troubleshooting Tools in Fabric Manager

The following Fabric Manager procedures are used to verify the VSAN, domain, FSPF, and zone configuration:

- Choose **Fabric*xx* > VSAN*xx*** to view the VSAN configuration in the Information pane.
- Choose **Fabric*xx* > VSAN*xx*** and select the **Host** or **Storage** tab in the Information pane to view the VSAN members.
- Choose **Fabric*xx* > VSAN*xx* > Domain Manager** to view the FCdomain configuration in the Information pane.
- Choose **Fabric*xx* > VSAN*xx* > FSPF** to view the FSPF configuration in the Information pane.
- Choose **Fabric*xx* > VSAN*xx* >** *zonesetname* to view the zone configuration for this VSAN. Zone configuration problems may appear to be a VSAN problem.

Common Troubleshooting Commands in the CLI

The following CLI commands are used to display VSAN, FCdomain, and FSPF information:

- **show vsan**
- **show vsan** *vsan-id*
- **show vsan membership**
- **show interface** *fc slot/port* **trunk** *vsan-id*
- **show** *vsan-id* **membership**
- **show vsan membership interface** *fc slot/port*
- **show fcdomain**
- **show fspf**
- **show fspf internal route vsan** *vsan-id*
- **show fcns database vsan** *vsan-id*

The following zone CLI commands may be useful to validate your configuration:

- **show zoneset name** *zonesetName vsan-id*
- **show zoneset active** *vsan-id*

Note An asterix (*) near the device listed by the **show zoneset active** CLI command indicates that the device is logged into the name server.

- **show zone** *vsan-id*
- **show zone status show** *vsan-id*

Note For more information on zoning issues, see Chapter 9, "Troubleshooting Zones and Zone Sets."

VSAN Issues

This section covers the following VSAN issues:

Host Cannot Communicate with Storage

Communication problems between a host and storage devices can be caused by port, VSAN, or zone issues.

Symptom Host cannot communicate with storage.

Table 7-1 Host Cannot Communicate with Storage

Symptom	Possible Cause	Solution
Host cannot communicate with storage.	Host and storage are not in the same VSAN.	Verify the VSAN membership. See the "Verifying VSAN Membership Using Fabric Manager" section on page 7-6 or the "Verifying VSAN Membership Using the CLI" section on page 7-6.
	xE port connecting to the remote switch is isolated.	See the "xE Port Is Isolated in a VSAN" section on page 7-7.
	Host and storage are not in the same zone.	See the "Zone and Zone Set Issues" section on page 9-4.

Verifying VSAN Membership Using Fabric Manager

To verify VSAN membership for host and storage devices using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > VSAN*xx*** and select the **Host** or **Storage** tab in the Information pane. Verify that both devices are in the same VSAN.

Step 2 If the host and storage are in different VSANs, verify which port is not in the correct VSAN and then follow these steps to change the port VSAN:

a. Highlight the host or storage in the Information pane. You see the link to that end device highlighted in blue in the map pane.

b. Right-click on the highlighted link and select **Interface Attributes** from the pop-up menu.

c. Set the PortVSAN field to the VSAN that holds the other end device and click **Apply Changes**.

Step 3 Right-click any ISL between the switches and select **Interface Attributes**. Select the **Trunk Config** tab and verify that the allowed VSAN list includes the VSAN found in Step 1.

Step 4 If the trunk is not configured for the VSAN, set the Allowed VSANs field to include the VSAN that the host and storage devices are on and click **Apply Changes.**

Verifying VSAN Membership Using the CLI

To verify VSAN membership for host and storage devices using the CLI, follow these steps:

Step 1 Use the **show vsan membership** command to see all the ports connected to your host and storage, and verify that both devices are in the same VSAN. Use this command on the switches that connect to your host or storage devices.

```
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
```

Step 2 If the host and storage are in different VSANs, use the **vsan database vsan** *vsan-id* **interface** CLI command to move the interface connected to the host and storage devices into the same VSAN.

Step 3 Use the **show interface** command to verify that the trunks connecting the end switches are configured to transport the VSAN found in Step 1.

```
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
```

Step 4 If the trunk is not configured for the VSAN, use the **interface** CLI command and then the **switchport trunk allowed vsan** CLI command in interface mode to add the VSAN to the allowed VSAN list for the interface that connects the host and storage devices.

xE Port Is Isolated in a VSAN

Symptom xE port is isolated in a VSAN.

Table 7-2 xE Port is Isolated in a VSAN

Symptom	Possible Cause	Solution
xE port is isolated in a VSAN.	E port connecting to the remote switch is isolated.	Verify the VSAN. See the "Resolving an Isolated E Port Using Fabric Manager" section on page 7-8 or the "Resolving an Isolated E Port Using Fabric Manager" section on page 7-8.
	TE port connecting to the remote switch is isolated.	See the "Resolving an Isolated ISL Using Fabric Manager" section on page 7-9 or the "Resolving an Isolated ISL Using the CLI" section on page 7-9
	Fabric timers misconfigured.	Use caution when changing fabric timers. See the "Resolving Fabric Timer Issues Using Fabric Manager" section on page 7-11 or the "Resolving Fabric Timer Issues Using the CLI" section on page 7-11.
	Port parameters misconfigured.	See the "Common Problems with Port Interfaces" section on page 6-10.
	Zoning mismatch.	See Chapter 9, "Troubleshooting Zones and Zone Sets."

Resolving an Isolated E Port Using Fabric Manager

To resolve VSAN isolation on an E port using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > FC Physical** and check the FailureCause column on the E port to verify that you have a VSAN mismatch problem.

Step 2 Choose **Switches > Interfaces > FC Physical** and set the PortVSAN field to correct a VSAN mismatch.

Resolving an Isolated E Port Using the CLI

To resolve VSAN isolation on an E port using the CLI, follow these steps:

Step 1 Use the **show interface** command to verify that the port is isolated because of a VSAN mismatch.

```
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 4
    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
```

Step 2 Use the **show vsan membership** CLI command to verify that the ports are in separate VSANs.

```
switch# show vsan membership
vsan 3 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 4 interfaces:
        fc2/5   fc2/13

vsan 4094(isolated_vsan) interfaces:
```

This sample output shows that all the interfaces on the switch belong to VSAN 3, with the exception of interface fc2/5 and fc2/13, which are part of VSAN 4.

Step 3 Use the **vsan database vsan** *vsan-id* **interface** CLI command to move the ports into the same VSAN.

Resolving an Isolated ISL Using Fabric Manager

Trunking E ports (TE ports) are similar to E ports except that 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 for ISL isolation on each VSAN.

To resolve VSAN isolation on a TE port using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > FC Physical** and check the FailureCause column on the TE port to verify that you have trunk problems.

Step 2 Choose **Switches > Interfaces > FC Physical** and select the **Trunk Failures** tab to determine the reason for the trunk problem.

Step 3 Correct the problem listed in the FailureCause column. See the "DPVM Config Database Not Activating" section on page 7-16 for domain misconfiguration problems. Choose **Switches > Interfaces > FC Physical** and set the PortVSAN field to to correct the VSAN misconfiguration problems.

Step 4 Repeat this procedure for all isolated VSANs on this TE port.

Resolving an Isolated ISL Using the CLI

Trunking E ports (TE ports) are similar to E ports except that 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 for ISL isolation on each VSAN.

To resolve VSAN isolation on a TE port using the CLI, follow these steps:

Step 1 Use the **show interface** command on the TE port to verify that you have an isolated VSAN.

```
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
```

The example shows the output of the **show interface** command with one or more isolated VSANs. Here, the TE port has one VSAN isolated.

Step 2 Use the **show interface fc** *slot/port* **trunk vsan** *vsan-id* command to verify the reason for VSAN isolation.

```
switch# show 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.

Step 3 Use the **show port internal info interface fc** *slot/port* command to determine the root cause of the VSAN isolation.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

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

fc2/14 - if_index: 0x0109C000, phy_port_index: 0x3c
  Admin Config - state(up), mode(TE), 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(0x3)
    description()
    Hw Capabilities: 0xb
    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 configured on
peer), fcid(0x000000)
    port init flag(0x0), current state [TE_FSM_ST_ISOLATED_VSAN_MISMATCH]
```

The last few lines of the command output provide a description of the reason for VSAN isolation for every isolated VSAN.

In this example, VSAN 7 is up, while two VSANs are isolated. VSAN 1 is isolated because of domain ID misconfiguration, and VSAN 8 is isolated because of VSAN misconfiguration.

Step 4 Correct the root cause. See the "DPVM Config Database Not Activating" section on page 7-16 for domain misconfiguration problems. Use the the **vsan** *vsan-id* **interface** CLI command to correct the VSAN misconfiguration problems.

Step 5 Repeat this procedure for all isolated VSANs on this TE port.

Resolving Fabric Timer Issues Using Fabric Manager

Use caution when changing fabric timers.

To resolve FC timer issues between VSANs using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > VSAN*xx* > VSAN Attributes** to verify that the fabric timers are inconsistent across the VSANs.

Step 2 Choose **Switches > FC Services > Timers and Policies.** You see the fabric timers in the Information pane.

Step 3 Click **Change Timeout Values** and set the timers and click **Apply**.

Resolving Fabric Timer Issues Using the CLI

Use caution when changing fabric timers.

To resolve fabric timer issues between VSANs using the CLI, follow these steps:

Step 1 Use the **show fctimer** CLI command to verify that the fabric timers are inconsistent across the VSANs.

Step 2 Use the **fctimer distribute** CLI command to enable CFS distribution for the fabric timers. Repeat this on all switches in this VSAN.

Step 3 Use the **fctimer** CLI command to set each timer.

Step 4 Use the **fctimer commit** command to save these changes and distribute them to all switches in the VSAN.

Troubleshooting Interop Mode Issues

To troubleshoot interop modes, refer to the switch to switch interop guide at the following website:

http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/mdsint/intgd.pdf

Dynamic Port VSAN Membership Issues

Dynamically assigning VSAN membership to ports is achieved by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature. DPVM offers flexibility and eliminates the need to reconfigure the VSAN to maintain fabric topology when a host or storage device connection is moved between two switches or between ports on the same switch. It retains the configured VSAN regardless of where a device is connected or moved.

Verify the following requirements when using DPVM:

- The interface through which the dynamic device connects to the Cisco MDS switch must be configured as an F port. FL ports do not support DPVM and no entries will be learned through an FL port.
- The static port VSAN of the F port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).

Note The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

Note If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

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

For more information on enabling DPVM, see one of the following guides:

- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Family Configuration Guide*

This section contains the following topics:

Troubleshooting DPVM Using Fabric Manager

To troubleshoot DPVM using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > All VSANs > DPVM** and select the **CFS** tab.

Step 2 Verify that the Oper and Global columns are enabled. If not, set the Admin drop-down menu to **enable** and the Global drop-down menu to **enable**. Then click **Apply Changes**.

Step 3 Select the **Actions** tab. Uncheck **AutoLearn Enable** if it is checked and click **Apply Changes**.

Step 4 Select the **Active Database** tab.

Step 5 Select **Pending** from the Compare To drop-down menu. You see a dialog box listing any differences between the active DPVM database and the pending database.

Step 6 Select the **CFS** tab and set Config Action to **commit** if there are any pending changes that you want to save. Click **Apply Changes.**

Step 7 Select the **Actions** tab and select activate from the Actions drop-down menu to activate the database. Click **Apply Changes**.

Troubleshooting DPVM Using the CLI

To troubleshoot DPVM using the CLI, follow these steps:

Step 1 Use the **show dpvm** CLI command in EXEC mode to verify that CFS distribution is enabled for DPVM. Optionally, use the **dpvm distribute** CLI command in config mode to enable CFS distribution if required.

Step 2 Use the **show dpvm status** CLI command in EXEC mode to verify that autolearning is disabled. Optionally, use the **no dpvm auto-learn** command in config mode if you need to disable autolearning before activating the database.

Step 3 Use the s**how dpvm pending-diff** CLI in EXEC mode command to compare the active and pending databases.
Optionally use the **dpvm commit** CLI command in config mode to commit any pending entries to the config database.

Step 4 Use the **dpvm activate** CLI in config mode command to activate the database.

DPVM Configuration Not Available

Symptom DPVM configuration is not available on Fabric Manager or CLI.

Table 7-3 DPVM Configuration not Available

Symptom	Possible Cause	Solution
DPVM configuration is not available on Fabric Manager or CLI.	DPVM has not been enabled.	DPVM must be enabled before it can be configured. Choose **Fabricxx > All VSANs > DPVM** and check the Status field in Fabric Manager or use the **show dpvm status** CLI command to verify that DPVM is not enabled. Set the Status field to **enable** in Fabric Manager and then click **Apply Changes** or use the **dpvm enable** CLI command to enable DPVM.

DPVM Database Not Distributed

Symptom DPVM databases are not distributed.

Table 7-4 DPVM Database not Distributed

Symptom	Possible Cause	Solution
DPVM databases are not distributed.	DPVM distribution has not been enabled on the local switch.	Choose **Fabricxx > All VSANs > DPVM** and select the CFS tab. Check the Global field in Fabric Manager or use the **show dpvm status** CLI command to verify that DPVM distribution is not enabled. Set the Global field to **enable** in Fabric Manager and then click **Apply Changes** or use the **dpvm distribute** CLI command to enable DPVM.
	DPVM distribution has not been enabled on one or more remote switches.	

DPVM Autolearn Not Working

The DPVM autolearn feature allows you to automatically populate the DPVM configuration database with all devices currently in the fabric. This feature is best used when you first turn on DPVM on a stable fabric. Once the devices are learned, you disable autolearning to populate the configuration database with these autolearned entries.

When you add a new device, it is best practices to manually add that device to the DPVM configuration database. If you turn on autolearning for a new device, you may add other devices that you did not intend to add.

Symptom DPVM autolearn does not work or is not getting enabled.

Table 7-5 DPVM Autolearn not Working

Symptom	Possible Cause	Solution
DPVM autolearn does not work or is not getting enabled.	DPVM active database may be absent.	Choose **Fabricxx > All VSANs > DPVM** and select the **Active Database** tab in Fabric Manager or use the **show dpvm database** CLI command to verify that DPVM is not enabled. Select the Actions tab and set the Action field to **activate** in Fabric Manager and then click **Apply Changes** or use the **dpvm activate** and **dpvm commit** CLI commands to create the DPVM active database.

Note When DPVM distribution is enabled, you must do an explicit commit for DPVM activate and autolearn to take effect.

No Autolearn Entries in Active Database.

Symptom There are no autolearn entries in the active database.

Table 7-6 No Autolearn Entries in Active Database.

Symptom	Possible Cause	Solution
There are no autolearn entries in the active database.	Autolearn is not enabled.	Choose **Fabricxx > All VSANs > DPVM** and select the **Actions** tab in Fabric Manager or use the **show dpvm status** CLI command to determine if autolearn is enabled. Check the **Auto Learn Enable** check box in Fabric Manager and click **Apply Changes** or use the **dpvm auto-learn enable** and **dpvm commit** CLI commands to enable autolearning.
	Port type is not supported.	Verify that the device you want to autolearn is connected to an F port. DPVM does not support FL, TE, FCIP, or PortChannels.

VSAN Membership not Added to Database.

Symptom The VSAN membership of the port is not added to the database.

Table 7-7 VSAN Membership not Added to Database.

Symptom	Possible Cause	Solution
The VSAN membership of the port is not added to the database.	Entry may be present in the config database.	Choose **Fabric*xx* > All VSANs > DPVM** and select the **Config Database** tab in Fabric Manager or use the **show dpvm database** CLI command to determine if the entry is present in the config database.
	DPVM distribution is enabled but a database change was not committed.	Choose **Fabric*xx* > All VSANs > DPVM** and select the **CFS** tab in Fabric Manager. Set the Config Action drop-down menu to **commit**. Or Use the **show dpvm pending** CLI command to determine if there are uncommitted changes. Use the **dpvm database** and **dpvm commit** CLI commands to commit any pending changes.

DPVM Config Database Not Activating

Symptom DPVM config database is not getting activated.

Table 7-8 DPVM Config Database not Activating

Symptom	Possible Cause	Solution
DPVM config database is not getting activated.	Conflicting entries may be present between the DPVM config and active databases.	Use the **dpvm database diff active conf** CLI command to determine if there are conflicting entries between the active and config databases. Choose **Fabricxx > All VSANs > DPVM** and select the **Actions t**ab in Fabric Manager. Set the Actions drop-down menu to **forceActivate** and Click **Apply Changes** or use the **dpvm activate force** and **dpvm commit** CLI commands to override the active database with the config database.

Cannot Copy Active to Config DPVM Database

Symptom Cannot copy the active DPVM database to the config database.

Table 7-9 DPVM Merge Failed

Symptom	Possible Cause	Solution
Cannot copy the active DPVM database to the config database.	Active database may be absent.	Choose **Fabricxx > All VSANs > DPVM** and select the **Active Database** tab in Fabric Manager or use the **show dpvm database** CLI command to verify that DPVM is not enabled. Select the **Actions** tab and set the Action field to **activate** in Fabric Manager and then click **Apply Changes** or use the **dpvm activate** and **dpvm commit** CLI commands to create the DPVM active database. Then copy the active database again.

Port Suspended or Disabled after DPVM Activation

Symptom A port in a static VSAN that was operational goes into suspend or disabled state after DPVM database activation.

Table 7-10 DPVM Merge Failed

Symptom	Possible Cause	Solution
A port in a static VSAN that was operational goes into suspend or disabled state after DPVM database activation.	DPVM database maps a connected device to a nonexistent VSAN.	Choose **Switches > Interfaces > FC Physical** in Fabric Manager or use the **show interface** CLI command to check the interface status for a dynamic VSAN related failure. Create the VSAN or map the device to another VSAN.

DPVM Merge Failed

Symptom DPVM merge failed.

Table 7-11 DPVM Merge Failed

Symptom	Possible Cause	Solution
DPVM merge failed.	DPVM operational parameters in the two merging fabrics are different.	Choose **Fabricxx > All VSANs > DPVM** and check the or use the **show dpvm** CLI command to verify the DPVM configuration in both fabrics. Manually reconcile any differences before attempting to merge the fabrics. Use the **show cfs merge status name dpvm** CLI command to show the merge status.

Domain Issues

This section includes the following topics:

Domain ID Conflict Troubleshooting

In a Fibre Channel network, the principal switch assigns 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 preexisting switches becomes the principal switch for the merged fabric.

The election of the new principal switch is characterized by the following rules:

- A switch with a populated 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.
- If both fabrics have a domain ID list, the priority between the two principal switches is determined by the configured switch priority. This is a user-settable parameter. The lower the value is, the higher the priority.
- If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the WWNs of the two switches. The lower value WWN has the higher priority.

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

- In Cisco SAN-OS Release 2.1(1a) and later releases, when connecting a single-switch fabric to a multi-switch fabric, a BF occurs and the switch with the better priority becomes the principal switch. In earlier releases, 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.
- In Cisco SAN-OS Release 2.1(1a) and later releases, when powering up a new switch in a multi-switch fabric, a BF occurs and the switch with the better priority becomes the principal switch. In earlier releases, when powering up a new switch in 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 a standalone switch, the new principal switch is determined by the administratively assigned priority if both switches are running Cisco SAN-OS Release 2.0(x) or earlier. If no priority is assigned (where the default priority is used in every switch), the principal switch is determined by the WWN. 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 assigned priority. If no priority is assigned (where the default value is used by every switch), the principal switch is determined by the WWN of the existing principal switches of the two fabrics.

Two switch fabrics might not merge. If two fabrics with two or more switches are connected, and they have at least one assigned domain ID in common, and the auto-reconfigure option is disabled (this option is disabled by default), then the E ports that are used to connect the two fabrics will be isolated due to domain ID overlap.

Switch Cannot See Other Switches in a VSAN

Symptom Switch cannot see other switches in a VSAN.

Table 7-12 Switch Cannot See Other Switches in a VSAN

Symptom	Possible Cause	Solution
Switch cannot see other switches in a VSAN.	Switch is isolated because of a domain ID overlap.	To resolve the problem, you can either change the overlapping static domain ID 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. See the "FC Domain ID Overlap" section on page 7-19.
	Fabric timers are misconfigured.	See the "Resolving Fabric Timer Issues Using Fabric Manager" section on page 7-11 or the "Resolving Fabric Timer Issues Using the CLI" section on page 7-11.

FC Domain ID Overlap

To resolve an FC domain ID overlap, you can either change the overlapping static domain ID 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.

- To assign a static domain ID, see the "Assigning a New Domain ID Using Fabric Manager" section on page 7-19 or the "Assigning a New Domain ID Using the CLI" section on page 7-20.
- To assign a dynamic domain ID after a fabric reconfiguration, see the "Using Fabric Reconfiguration for Domain ID Assignments" section on page 7-21.

You may see the following system message in the message log when a domain ID overlap occurs:

Error Message `PORT-5-IF_DOWN_DOMAIN_OVERLAP_ISOLATION: Interface [chars] is down (Isolation due to domain overlap).`

Explanation The interface is isolated because of a domain overlap.

Recommended Action Use the **show fcdomain domain-lis**t to determine which domain IDs are overlapping. Us the fcdomain Use the **fcdomain domain** *domain-id* **[static | preferred] vsan** *vsan-id* CLI command or similar Fabric Manager procedure to change the domain ID for one of the overlapping domain IDs.

Assigning a New Domain ID Using Fabric Manager

All devices attached to the switch in the VSAN get a new FC ID when a new domain ID is assigned. Some hosts or storage devices may not function as expected if the FC ID of the host or storage device changes.

To verify FC domain ID overlap and reassign a new Domain ID using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > FC Physical** and check the FailureCause column for an isolation or domain overlap status.

Step 2 Choose **Fabric*xx* > VSAN*xx* > Domain Manger** to view which domains are currently in the VSAN.

Step 3 Repeat Step 2 on the other switch to determine which domain IDs overlap.

Step 4 Select the **Configuration** tab and set Config Domain and Config Type to change the domain ID for one of the overlapping domain IDs.

- The static option tells the switch to request that particular domain ID. If it does not get that particular address, it will isolate itself from the fabric.
- The preferred option has the switch request a specified domain ID. If that ID is unavailable, it will accept another ID.

Step 5 Set the Restart drop-down menu to **disruptive** and click **Apply Changes** to restart the Domain Manager.

Note While the static option can be applied to runtime after a disruptive or nondisruptive restart, the preferred option is applied to runtime only after a disruptive restart.

Assigning a New Domain ID Using the CLI

All devices attached to the switch in the VSAN get a new FC ID when a new domain ID is assigned. Some hosts or storage devices may not function as expected if the FC ID of the host or storage device changes.

To verify FC domain ID overlap and reassign a new Domain ID using the CLI, follow these steps:

Step 1 Issue the **show interface** command. The following example output shows the isolation error message.

```
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
```

Step 2 Use the **show fcdomain domain-list vsan** *vsan-id* command to view which domains are currently in your fabric.

```
switch1# show fcdomain domain-list vsan 2

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]
---------    -----------------------
```

Step 3 Repeat Step 2 on the other switch to determine which domain IDs overlap.

```
switch2# show fcdomain domain-list vsan 2

Number of domains: 1
Domain ID              WWN
---------    -----------------------
0x4b(75)    20:01:00:05:30:00:13:9e [Local][Principal]
---------    -----------------------
```

In this example, switch 2 is isolated because of a domain ID 75 overlap.

Step 4 Use the **fcdomain domain** *domain-id* **[static | preferred] vsan** *vsan-id* CLI command to change the domain ID for one of the overlapping domain IDs.

- The static option tells the switch to request that particular domain ID. If it does not get that particular address, it will isolate itself from the fabric.
- The preferred option has the switch request a specified domain ID. If that ID is unavailable, it will accept another ID.

Step 5 Use the **fcdomain restart disruptive vsan** CLI command to restart the Domain Manager.

Note While the static option can be applied to runtime after a disruptive or nondisruptive restart, the preferred option is applied to runtime only after a disruptive restart.

Using Fabric Reconfiguration for Domain ID Assignments

You can use a fabric reconfiguration to reassign domain IDs and resolve any overlapping domain IDs. 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 a new domain IDs to be assigned to the different switches.

Caution A disruptive reconfiguration might affect data traffic.

Using Fabric Reconfiguration for Domain ID Assignments with Fabric Manager

To use fabric reconfiguration to reassign domain IDs for a particular VSAN using Fabric Manager, follow these steps:

Step 1 Choose **Switches > Interfaces > FC Physical** and select the **Domain Manager** tab in the Information pane.

Step 2 Uncheck the **RcfReject** check box and click **Apply Changes** to disable RCF rejection.

Step 3 Choose **Fabric*xx* > VSAN*xx* > Domain Manager** in the Logical Domain pane.

Step 4 Click the **Configuration** tab in the Information pane and set the Config Type drop-down menu to **preferred** to remove any static domain ID assignments.

Step 5 Check the **AutoReconfigure** check box to enable the auto-reconfiguration option.

Step 6 Set the Restart drop-down menu to **disruptive** and click **Apply Changes** to restart the Domain Manager.

Using Fabric Reconfiguration for Domain ID Assignments with the CLI

To use fabric reconfiguration to reassign domain IDs for a particular VSAN using the CLI, follow these steps:

Step 1 Use the **show fcdomain domain-list** CLI command to determine if you have statically assigned domain IDs on the switches.

Step 2 If you have statically assigned domain IDs, use the **no fcdomain domain** CLI command to remove the static assignments.

Step 3 Use the **show fcdomain vsan** CLI command to determine if you have rcf-reject option enabled.

```
switch# show fcdomain vsan 1
The local switch is a Subordinated Switch

Local switch run time information:
      State: Stable
      Local switch WWN:    20:01:00:05:30:00:51:1f
      Running fabric name: 10:00:00:60:69:22:32:91
      Running priority: 128
      Current domain ID: 0x64(100) ß verify domain id

Local switch configuration information:
        State: Enabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 41:6e:64:69:61:6d:6f:21
        Configured priority: 128
        Configured domain ID: 0x64(100) (preferred)

Principal switch run time information:
       Running priority: 2

Interface              Role          RCF-reject
----------------    -------------    ------------
fc2/1               Downstream       Enabled
fc2/2               Downstream       Disabled
fc2/7               Upstream         Disabled
----------------    -------------    ------------
```

Step 4 If you have the rcf-reject option enabled, use the **interface** CLI command and then the **no fcdomain rcf-reject vsan** CLI command in interface mode.

Step 5 Use the **fcdomain auto-reconfigure vsan** CLI command in the EXEC mode on both switches to enable auto-reconfiguration after a Domain Manager restart.

Step 6 Use the **fcdomain restart disruptive vsan** CLI command to restart the Domain Manager.

FSPF Issues

The implementation of VSANs dictates that each configured VSAN support a separate set of fabric services. One such service is the FSPF routing protocol, which can be independently configured per VSAN. Therefore, within each VSAN topology, FSPF can be configured to provide a unique routing configuration and resulting traffic flow. Using the traffic engineering capabilities offered by VSANs allows a greater control over traffic within the fabric and a higher utilization of the deployed fabric resources.

This section describes how to identify and resolve Fabric Shortest Path First (FSFP) problems. It includes the following topics:

Troubleshooting FSPF

Figure 7-1 shows a single-VSAN topology.

Figure 7-1 Single VSAN Topology

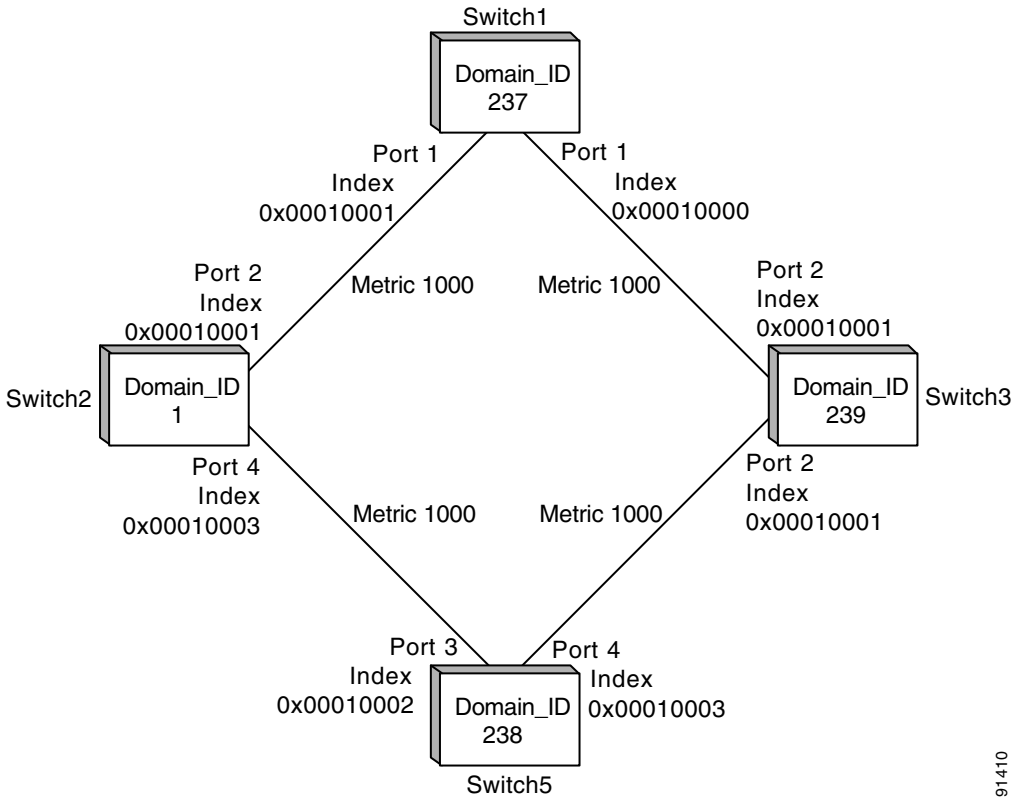

For the purpose of this example, assume that all interfaces are located in VSAN 1.

Troubleshooting FSPF Using Device Manager

To troubleshoot FSPF using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF** and select the **LSDB LSRs** tab to verify the link state records in the FSPF database.

- The VSANId/ DomainId column shows the domain's view of the fabric topology.
- The AdvDomainId column shows which domain is the owner of the LSR (link state record).
- The Age value is a 16-bit counter starting at 0x0000, incremented by one for each switch during flooding and by one for each second held in the database. This field is used as a tie-breaker if Incarnation numbers are the same.
- The IncarnationNumber is a 32-bit value between 0x80000001 and 0x7FFFFFFF that is incremented by one each time the originating switch transmits an LSR. This is used first before the Age value.

Step 2 Choose **FC > Advanced > FSPF** and select the **LSDB Links** tab to verify that each path is in the FSPF database.

Step 3 Choose **FC > Advanced > FSPF** and select the **Interfaces** tab to verify that the FSPF parameters are correct for each interface and verify that the AdminStatus is up.

- The Cost column shows the cost of the path out of the interface.
- The Intervals column shows the configured FSPF timers for this interface, which must match on both sides.
- The State column shows the full or adjacent state if the interface has sent and received all database exchanges and required Acks. The port is now ready to route frames.
- The Neighbors column shows FSPF neighbor information.

Step 4 Choose **FC > Advanced > FSPF** and select the **Statistics** or **InterfaceStats** tab to verify that there are no excessive errors present.

Troubleshooting FSPF Using the CLI

To troubleshoot FSPF using the CLI, follow these steps:

Step 1 Use the **show fspf database vsan** CLI command to verify that each path is in the FSPF database.

```
switch1# show fspf database
FSPF Link State Database for VSAN 2 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

FSPF Link State Database for VSAN 2 Domain 237 <-----------LSR for another switch
LSR Type                = 1
Advertising domain ID   = 237 -----7
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

FSPF Link State Database for VSAN 2 Domain 238  <-----------LSR for another switch
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

FSPF Link State Database for VSAN 2 Domain 239 <-----------LSR for another switch
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. The domain 1 view of the fabric topology.
2. 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. The path to domain 237, switch 1.
6. The path to domain 238, switch 5.
7. Switch 1, domain ID 237 is the owner.
8. The path to domain 239, switch 3.
9. The path to domain 1, switch 2.

Step 2 Use the **show fspf vsan** *vsan-id* **interface** CLI command to verify that the FSPF parameters are correct for each interface and verify that the interface is in the FSPF active state.

```
switch1# show fspf vsan 2 interface fc1/2
FSPF interface fc1/2 in VSAN 2
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.

10. The cost of the path out this interface.
11. The configured FSPF timers for this interface, which must match on both sides.
12. Either Full State or Adjacent. Sent and received all database exchanges and required Acks. Port is now ready to route frames.
13. FSPF neighbor information.

Step 3 Use the **show fspf internal route vsan** CLI command to verify that all Fibre Channel routes are available.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
switch1# show fspf internal route vsan 2
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.

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.

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.

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

Loss of Two-Way Communication

If FSPF is misconfigured, then the switches will not reach the “two-way” state.

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

- The port enters Init state and removes its neighbor’s domain ID from the Recipient Domain ID field and inserts 0xFFFFFFFF.
- FSPF removes the Inter-Switch Link (ISL) from the topology database.
- New link state records (LSRs) are flooded to adjacent switches to notify them that the FSPF database has changed.

Symptom Traffic is not being routed through the fabric.

Table 7-13 Traffic is not Being Routed Through the Fabric

Symptom	Possible Cause	Solution
Traffic is not being routed through the fabric.	FSPF hello interval misconfigured.	See the "Resolving a Wrong Hello Interval on an ISL Using Device Manager" section on page 7-28 or the "Resolving a Wrong Hello Interval on an ISL Using the CLI" section on page 7-29.
	FSPF retransmit time misconfigured.	See the "Resolving a Mismatched Retransmit Interval on an ISL Using Device Manager" section on page 7-30 or the "Resolving a Mismatched Retransmit Interval on an ISL Using the CLI" section on page 7-30.
	FSPF dead interval misconfigured.	See the "Resolving a Mismatch in Dead Intervals on an ISL Using Fabric Manager" section on page 7-31 or the "Resolving a Mismatch in Dead Intervals on an ISL Using the CLI" section on page 7-31.
	There is a region mismatch on the switch.	See the "Resolving a Region Mismatch Using Fabric Manager" section on page 7-32 or the "Resolving a Region Mismatch Using the CLI" section on page 7-32.

Resolving a Wrong Hello Interval on an ISL Using Device Manager

To resolve a wrong hello interval on an ISL using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF** and select the **Interfaces** tab to verify that the FSPF parameters are correct for each interface and check the Hello interval column and the State column.

- The Intervals column shows the configured FSPF timers for this interface, which must match on both sides.
- The State column shows the full or adjacent state if the interface has sent and received all database exchanges and required Acks. The port is now ready to route frames.

Step 2 Repeat Step 1 to determine the value of the hello interval on the adjacent switch.

Step 3 Fill in the Hello field to change the hello interval and click **Apply**.

Resolving a Wrong Hello Interval on an ISL Using the CLI

To resolve a wrong hello interval on an ISL using the CLI, follow these steps:

Step 1 Use the **debug fspf all** CLI command and look for wrong hello interval messages.

```
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
```

Tip We recommend that you open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 2 Use the **undebug all** command to turn off debugging.

Step 3 Use the **show fspf internal route vsan** to show FSPF information.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
switch1# show fspf internal route vsan 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. There is no second path to domain 238, through domain 1 switch 2.
2. There is no direct path to domain 1 switch 2; traffic must travel through three ISLs. This is based on the route cost column.

Step 4 Use the **show fspf vsan** *vsan-id* **interface** CLI command to view the FSFP configuration.

```
switch1# show fspf vsan 1 interface 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
```

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

Step 5 Repeat Step 4 to determine the value of the Hello timer on the adjacent switch.

```
switch2# show fspf v 1 interface fc2/16
FSPF interface fc2/16 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. The neighbor FSPF Hello interval is set to the default (20 seconds).
2. FSPF is not in full state, indicating a problem.

Step 6 Use the **interface** CLI command and then the **fspf hello-interval** CLI command in interface mode to change the default Hello interval.

Resolving a Mismatched Retransmit Interval on an ISL Using Device Manager

To resolve a mismatched retransmit interval on an ISL using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF** and select the **Interfaces** tab to verify that the FSPF parameters are correct for each interface and check the Retransmit interval column and the State column.

- The Intervals column shows the configured FSPF timers for this interface, which must match on both sides.
- The State column shows the full or adjacent state if the interface has sent and received all database exchanges and required Acks. The port is now ready to route frames.

Step 2 Repeat Step 1 to determine the value of the retransmit interval on the adjacent switch.

Step 3 Fill in the Retransmit field to change the retransmit interval and click **Apply**.

Resolving a Mismatched Retransmit Interval on an ISL Using the CLI

To resolve a mismatched retransmit interval on an ISL using the CLI, follow these steps:

Step 1 Use the **show fspf vsan** *vsan-id* **interface** CLI command to view the FSFP configuration.

```
switch1# show fspf vsan 1 interface 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 10 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. The retransmit interval is not set to the default, so you should check the neighbor configuration to make sure it matches.
2. FSPF is not in FULL state, indicating a problem.

Step 2 Repeat Step 1 to determine the value of the retransmit interval on the adjacent switch.

```
switch2# show fspf v 1 interface fc2/16
FSPF interface fc2/16 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. The neighbor retransmit interval interval is set to the default (5 seconds).
2. FSPF is not in FULL state, indicating a problem.

Step 3 Use the **interface** CLI command and then the **fspf retransmit-interval** CLI command in interface mode to change the retransmit interval.

Resolving a Mismatch in Dead Intervals on an ISL Using Fabric Manager

To resolve a mismatch of dead intervals on an ISL using Fabric Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF** and select the **Interfaces** tab to verify that the FSPF parameters are correct for each interface and check the Dead interval column and the State column.

- The Intervals column shows the configured FSPF timers for this interface, which must match on both sides.
- The State column shows the full or adjacent state if the interface has sent and received all database exchanges and required Acks. The port is now ready to route frames.

Step 2 Repeat Step 1 to determine the value of the dead interval on the adjacent switch.

Step 3 Fill in the Dead field to change the dead interval and click **Apply**.

Resolving a Mismatch in Dead Intervals on an ISL Using the CLI

To identify a mismatch in dead intervals on an ISL, follow these steps:

Step 1 Use the **debug fspf all** CLI command and look for wrong dead interval messages.

```
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
```

Tip We recommend that you open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 2 Use the **undebug all** command to turn off debugging.

Step 3 Use the **show fspf vsan** <*vsan-id*> **interface** to show FSPF information.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
switch1# show fspf vsan 1 interface 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 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. The dead timer is not set to the default, so you should check the neighbor configuration.
2. FSPF is not in full state, which indicates a problem.

Step 4 Use the **interface** CLI comma nd and then the **fspf dead-interval** CLI command in interface mode to change the dead interval.

Resolving a Region Mismatch Using Fabric Manager

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

Step 1 Choose **FC > Advanced > FSPF** and select the **General** tab to verify the RegionId.

Step 2 Repeat Step 1 to determine the value of the region on the adjacent switch.

Step 3 Fill in the RegionId field to change the region and click **Apply**.

Resolving a Region Mismatch Using the CLI

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

Step 1 Use the **show fspf vsan** CLI command to display the currently configured region in a VSAN.

```
switch# show fspf vsan 99

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 */
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 Use the **debug fspf all** CLI command and look for nonexistent region messages.

```
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 fc1/1 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT
Jan  5 00:39:45 fspf: Interface fc1/2 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT -----2
```

1. The neighbor switch advertising region is 0.
2. FSPF is in init state for each ISL.

Tip We recommend that you open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 3 Use the **undebug all** command to turn off debugging.

Step 4 Use the **show fspf** CLI command to show FSPF configuration and check the autonomous region.

Step 5 Use the **fspf config vsan** CLI command to enter the fspf configuration mode and use the **region** CLI command to change the region.

The region must match on all switches in the VSAN.

CHAPTER 8

Troubleshooting IVR

This chapter describes how to troubleshoot and resolve inter-VSAN routing (IVR) configuration issues in the Cisco MDS 9000 Family of multilayer directors and fabric switches. It includes the following sections:

Overview

Troubleshooting IVR involves checking the configuration of domain IDs, VSANs, border switches, and zone sets. Configuration problems with IVR can prevent devices from communicating properly.

Prior to Cisco MDS SAN-OS Release 2.1(1a), IVR required unique domain IDs for all switches in the fabric. As of Cisco MDS SAN-OS Release 2.1(1a), you can enable IVR Network Address Translation (NAT) to allow non-unique domain IDs. This feature simplifies the deployment of IVR in an existing fabric where non-unique domain IDs might be present.

Note By default, IVR-NAT is not enabled.

Best Practices

This section provides the best practices for implementing IVR:

- Use Fabric Manager to configure IVR.

  Using Fabric Manager to configure IVR can help avoid errors and will ensure that the same IVR configuration is applied to all IVR enabled switches.

- Use IVR-NAT. If you do not use IVR-NAT, you must use non-overlapping domains across VSANs associated with IVR.

Note If you are using IVR-NAT, you are not required to use non-overlapping domains across VSANs.

- For large installations, do not spread IVR zone members across many switches.

  The VSAN rewrite table is limited to 4096 entries, and the entries are per-domain, not per-end device, so it is best to minimize the number of switches that contain IVR zone members in very large implementations.
- Use static domain IDs. This prevents changes in domain IDs that may conflict with virtual domain ID assignments.
- Allow for multiple paths between the IVR zone members. Implement redundant path designs whenever possible.
- Set the default zone policy to deny and avoid using the **force** option when activating the IVR zone set.

  In normal Fibre Channel environments, it is generally considered a best practice to set the default zone policy to deny. Because members of IVR zones cannot exist in the default zone, activation of an IVR zone set using the **force** option may lead to traffic disruption if IVR zone members previously existed in a default zone policy of permit.
- Use IVR auto-topology. If you do not use IVR auto-topology, use CFS distribution to ensure that the same IVR topology is applied to all IVR-enabled switches.
- Configure IVR only in the relevant border switches.
- Configure IVR-enabled VSANs in no interop (default) mode or interop 1 mode.
- Turn RDI mode on. This ensures that the switch will not assign used domain IDs and is compatible with third-party switches. In Cisco SAN-OS Release 2.0(x) and earlier, existing domain IDs are reserved in a local database. In Cisco SAN-OS Release 2.1(1a) and later, domain IDs are dynamically reserved using RDI.

Note Contact your customer support representative for more information regarding this feature (specifically for CSCei88345 and Field Notice 62187).

Transit VSANs

Follow these guidelines when configuring transit VSANs:

- Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though not prohibited) to provide connectivity.
  - If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.
- Traffic between the edge VSANs traverses only the shortest IVR path.
- Transit VSAN information is common to all IVR zones. Sometimes a transit VSAN can also be an edge VSAN in another IVR zone.

Border switches

Always follow these guidelines when configuring border switches:

- Border switches require Cisco SAN-OS Release 1.3(1) or higher.
- A border switch must be a member of two or more VSANs.
- A border switch that facilities IVR communications must be IVR enabled.
- For redundant paths between active IVR zone members, IVR can (optionally) be enabled on additional border switches.
- The VSAN topology configuration must be updated before a border switch is added or removed.

Initial Troubleshooting Checklist

Begin troubleshooting IVR issues by checking the following issues first:

Checklist	Checkoff
Verify that IVR is enabled on all border switches involved in IVR.	❑
Verify that you have the correct license installed (SAN_EXTENSION for IVR over FCIP or ENTERPRISE_PKG for IVR over Fibre Channel).	❑
Verify that the IVR configuration is the same on all IVR-enabled switches.	❑
Verify that the IVR zone is part of the active IVR zone set.	❑
Verify that you have an active zone set or that you activate the IVR zone set using the **force** option.	❑
Verify that you have added IVR virtual domains to the allowed domain ID list if you have a Cisco SN5428 storage router or a Cisco MDS 9020 switch in your fabric.	❑

If you change any FSPF link cost, ensure that the FSPF path cost (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.

This section includes the following topics:

Verifying IVR Configuration Using Fabric Manager

To verify your IVR configuration using Fabric Manager, follow these steps:

Step 1 Choose **Fabricxx > All VSANs > IVR** to verify your IVR configuration.

Step 2 Select the **CFS** tab to verify that the Oper column is enabled and the Global column is enabled for CFS distribution. Check the LastResult column for the status of the last CFS action.

Step 3 Select the **Action** tab to determine if auto topology and IVR NAT are enabled.

Step 4 Select the **Local Topology** and **Active Topology** tabs to verify your IVR VSAN topology.

Step 5 Choose **Fabricxx > All VSANs > Domain Manager** to verify unique domain IDs if IVR NAT is not enabled.

Step 6 Choose **Zone > IVR > Edit Local Full Zone Database** to verify your IVR zones and zone sets and to verify that you have activated your IVR zone set. The active IVR zone set name appears in bold.

Verifying IVR Configuration Using the CLI

Several commands involving multiple configuration tasks can be used to verify the IVR configuration.

Table 8-1 *CLI Commands for Verification of IVR*

CLI Command	Description
show fcdomain domain-list	Verifies unique domain ID assignment. If a domain overlap exists, edit and verify the allowed-domains list or manually configure static, non-overlapping domains for each participating switch and VSAN.
show interface brief	Verifies if the ports are operational, VSAN membership, and other configuration settings covered previously.
show fcns database	Verifies the name server registration for all devices participating in the IVR.
show zoneset active	Displays zones in the active zone set. This should include configured IVR zones.
show ivr fcdomain	Displays the IVR persistent fcdomain database.
show ivr internal	Shows the IVR internal troubleshooting information.
show ivr pending-diff	Shows the IVR pending configuration.
show ivr service-group	Shows the difference between the IVR pending and configured databases.
show ivr tech-support	shows information that is used by your customer support representative to troubleshoot IVR issues.
show ivr virtual-domains	Shows IVR virtual domains for all local VSANs.
show ivr virtual-fcdomain-add-status	Shows IVR virtual fcdomain status.
show ivr vsan-topology	Verifies the configured IVR topology.
show ivr zoneset	Verifies the IVR zone set configuration.
show ivr zone	Verifies the IVR zone configuration.
clear ivr zone database	Clears all configured IVR zone information. **Note** Clearing a zone set erases only the configured zone database, not the active zone database.

The following **show internal** commands can be useful for troubleshooting IVR issues.

```
add-rw                 Show ivr fcid rewrite fsm internals
  adv_vsans              Show IVR advertise VSANs for a native VSAN and domain
  area-port-allocation   Show IVR area-port allocation
  capability-fsm         Show IVR capability fsm internal debug information
  commit-rw              Show ivr fcid rewrite fsm internals
  debug-log-buffer1      Show IVR debug-log buffer
  del-rw                 Show ivr fcid rewrite fsm internals
  dep                    Show ivr dep internals
  device-list            Show ivr device list
  distribution           Show ivr distribution internals
  domain-capture-list    Show ivr domain controller capture list
  drav-fsm               Show DRAV FSM details
  event-history          Show ivr internal event history
  fcid-rewrite-fsm       Show ivr fcid rewrite fsm internals
  fcid-rewrite-list      Show ivr fcid rewrite entries
  fsmtca                 Show IVR FSM transition statistics
  global-data            Show ivr global data
  mem-stats              Show memory statistics
  nhvsan-change          Show ivr fcid rewrite fsm internals
  plogi-captured-list    Show ivr PLOGI captured
  pnat                   Show IVR payload NAT internal information
  pvm                    Show IVR PV Master internal information
  tu-fsm                 Show TU FSM internal debug information
  vdri-fsm               Show VDRI FSM internal debug information
  virtual-domains        Show IVR capability fsm internal debug information
  vsan-rewrite-list      Show ivr vsan rewrite list
  vsan-topology          Show internal information on IVR VSAN topology
  vsan-topology-graph    Show IVR VSAN Topology graph internal debug information
  zone-fsm               Show ivr zone fsm internals
```

Limitations and Restrictions

Limit the use of IVR NAT with write acceleration. Enabling IVR NAT on the same switch where write acceleration is enabled over a Port Channel of multiple FCIP links might result in frames from the source to the destination not transferring.

Design your SAN to properly use IVR and IVR zones. Design IVR zones to enable communications between devices that require it. Do not group all devices into one IVR zone if you do not require all those devices to communicate with each other.

Table 8-1 shows the limitations to the IVR configuration based on the Cisco SAN-OS release.

Table 8-1 IVR Configuration Limitations

Parameter per Fabric	Cisco SAN-OS 2.0(1b)	Cisco SAN-OS 2.1(1a) or later
IVR zone members	2000	10000
IVR zones	200	2000
IVR zone sets	32	32
VSANs	64	128
IVR-enabled switches	128	128

Note Two VSANS with the same VSAN ID combined with a unique AFID count as two VSANs in the total number of allowed VSANs per fabric.

IVR Enhancements by Cisco SAN-OS Release

Table 8-2 lists the IVR enhancements by Cisco SAN-OS release.

Table 8-2 IVR Enhancements by Cisco SAN-OS Release

Cisco SAN-OS Release	IVR Enhancement
Release 2.1(2)	Persistent FC IDs and domains for IVR
Release 2.1(1a)	• IVR NAT • AFIDs • Auto-topology • Virtual domains added to remote domain lists • IVR LUN zoning • IVR QoS zoning • Service group
Release 2.0(1)	IVR with CFS support
Release 1.3(4a)	Virtual domains added to remote domain lists.
Release 1.3(1)	IVR introduced.

Common IVR Problems

This section describes the problems associated with IVR. This section includes the following topics:

IVR allows device discovery across VSANs. IVR also supports FC ping and FC traceroute across VSANs using the following criteria:

- Either FC ID or pWWN can be used.
- Must be initiated from a switch with an active IVR zone member.

IVR Licensing Issues

To use IVR, you must obtain the correct licenses for the IVR features you are using and install those licenses on every IVR-enabled switch in your fabric. Table 8-3 shows which license to purchase, based on the IVR feature you are using and the module or chassis you have enabled IVR on.

Table 8-3 License Requirements for IVR

IVR Feature	Chassis or Module Type	License Required	Number of Licenses
IVR over Fibre Channel	All	ENTERPRISE_PKG	One per IVR-enabled chassis
IVR over FCIP	MDS 9216i [1]	None	None
	MPS-14/2	SAN_EXTN_OVER_IPS2	One per module running IVR Over FCIP
	IPS-8	SAN_EXTN_OVER_IP	
	IPS-4	SAN_EXTN_OVER_IPS4	

1. Cisco MDS 9216i enables the SAN_EXTENSION features without a license for the two Gigabit Ethernet ports on the integrated supervisor card.

Note If you are using IVR over FCIP and Fibre Channel, you need the ENTERPRISE_PKG as well as the appropriate SAN extension license as shown in Table 8-3.

Tip Be sure to enter the correct chassis serial number when purchasing your license packages. Choose **Switches > Hardware** and check the SerialNo Primary for the switch chassis in Fabric Manager or use the **show license host-id** CLI command to obtain the chassis serial number for each switch that requires a license. Your license will not operate if the serial number used does not match the serial number of the chassis you are installing the license on.

See Chapter 4, “Troubleshooting Licensing,” for complete details on troubleshooting licensing issues.

Cannot Enable IVR

Symptom Cannot enable IVR.

Table 8-4 Cannot Enable IVR

Symptom	Possible Cause	Solution
Cannot enable IVR.	License not installed and grace period has expired.	Purchase and install the appropriate licenses. See the "IVR Licensing Issues" section on page 8-7.
	Switch not running Cisco SAN-OS Release 1.3(1) or later.	Upgrade to the Cisco SAN-OS release required for the IVR features you want to use. See Table 8-1 and Chapter 2, "Troubleshooting Installs, Upgrades, and Reboots."
	Using IVR auto topology but CFS distribution is not enabled.	Choose **Fabricxx > All VSANs > IVR** and set the Global drop-down menu to **enable**. Click **Apply Changes**. Or use the **ivr distribute** CLI command before enabling IVR.

IVR Network Address Translation Fails

Symptom IVR NAT fails.

Table 8-5 IVR NAT Fails

Symptom	Possible Cause	Solution
IVR NAT fails.	Internal message payload uses destination ID.	IVR NAT modifies the destination ID in the Fibre Channel header. If this same destination ID appears inside the message payload, Cisco SAN-OS may not detect it and IVR NAT fails. Disable IVR NAT and ensure that all domain IDs are unique. Refer to the *Cisco MDS 9000 Family Configuration Guide* at the following website for a list of payloads that work with IVR NAT when the payload includes the destination ID: http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/2_0/cliguide/part_4/ivr.htm#wp1176738
	Some switches are running IVR without NAT.	You cannot combine IVR and IVR NAT in the same VSAN. Use the same IVR configuration on all switches. Deactivate the active zone set before converting to IVR or IVR NAT.

IVR Zone Set Activation Fails

Symptom IVR zone set activation fails.

Table 8-6 IVR Activation Fails

Symptom	Possible Cause	Solution
IVR zone set activation fails.	Overlapping domain IDs.	Use static domain IDs to assign unique domain IDs to each switch in the VSAN or use IVR NAT. Choose **Fabricxx > All VSANs > Domain Manager** in Fabric Manager or use the **fcdomain domain** *domain-id* **[static | preferred] vsan** *vsan-id* CLI command
	Default zone policy is permit.	Choose **Zone > IVR > Edit Local Full Zone Database** in Fabric Manager. Right-click the IVR zone set that you want to activate and select **Activate**. Check the **Create Active Zone Set if none Present** check box or use the **force** option with the **ivr zoneset activate** CLI command.
	Default zone policy is deny and no active zone set present.	
	No active zone set.	No zone set has been activated. See the "Troubleshooting Zone Set Activation" section on page 9-8 to activate a zone set on an IVR-enabled switch, or use the **force** option when activating the IVR zone set.

Border Switch Fails

If an IVR-enabled switch fails, you must update the IVR topology to reflect this change if you are not using auto topology.

Symptom Border switch fails.

Table 8-7 Border Switch Fails

Symptom	Possible Causes	Solutions
Border switch fails.	IVR topology incorrect.	Choose **Fabricxx > All VSANs > IVR** and select the **Action** tab in Fabric Manager. Check the **Auto Discover Topology** check box and click **Apply Changes.** Select the **CFS** tab and set ConfigAction to **commit** and click **Apply Changes**. Or use the **ivr vsan topology auto** CLI command to automatically reconfigure the IVR topology, or use the **ivr vsan topology database** CLI command to manually reconfigure the IVR topology.

Traffic Does Not Traverse IVR Path

Symptom Traffic does not traverse the IVR path.

Table 8-8 Traffic Does Not Traverse IVR Path

Symptom	Possible Cause	Solution
Traffic does not traverse the IVR path.	Fabric includes an SN5428 or MDS 9020 switch and you have not added the IVR virtual domains to the remote VSAN domain lists.	Choose **Fabricxx > All VSANs > IVR** and select the **Action** tab in Fabric Manager.Fill in the Create Virtual Domains for VSAN field and click **Apply Changes.** Select the **CFS** tab and set ConfigAction to **commit** and click **Apply Changes.** Or use the i**vr virtual-fcdomain-add vsan-ranges** CLI command to add existing and future virtual domains to the domain list for the selected VSANs. Repeat this on all edge VSANs.
	Internal message payload uses destination ID.	See the "IVR Network Address Translation Fails" section on page 8-8.

Link Isolated

Symptom Link isolated.

Table 8-9 Link Isolated

Symptom	Possible Cause	Solution
Link isolated.	Virtual domain overlap.	Choose **Fabricxx > All VSANs > Domain Manager** in Fabric Manager to verify a domain overlap. Choose **Fabricxx > All VSANs > IVR** and select the **Action** tab in Fabric Manager.Fill in the Create Virtual Domains for VSAN field and click **Apply Changes.** Select the **CFS** tab and set ConfigAction to **commit** and click **Apply Changes.** Or use the **show fcdomain domain-list** CLI command to verify a domain overlap. Use the **ivr widthdraw domain** CLI command to remove the overlapped domain. Use persistent FC IDs to reassign the overlapped domain. Use the i**vr virtual-fcdomain-add vsan-ranges** CLI command to add existing and future virtual domains to the domain list for the selected VSANs. Repeat this on all edge VSANs.
	Internal message payload uses destination ID.	See the "IVR Network Address Translation Fails" section on page 8-8.

Persistent FC ID for IVR Failed

Symptom Persistent FC ID for IVR failed.

Table 8-10 Persistent FC ID for IVR Failed

Symptom	Possible Cause	Solution
Persistent FC ID for IVR failed.	Selected virtual FC ID does not match the assigned virtual domain.	Use the **show ivr fcdomain database** CLI command to verify the virtual domain ID. Use the **native-autonomous-fabric-num** CLI command to assign the virtual domain and then use the **pwwn** CLI command to map the pWWN to an appropriate FC ID that matches the virtual domain ID. Refer to the *Cisco MDS 9000 Family Configuration Guide* for the related procedure to configure Persistent FC IDs for IVR.

LUN Configuration Failure in IVR Zoning

Symptom LUN configuration failed in IVR zoning.

Table 8-11 LUN Configuration Failure in IVR Zoning

Symptom	Possible Cause	Solution
LUN configuration failed in IVR zoning.	One or more switches in the VSAN are not running Cisco MDS SAN-OS Release 2.1(1a) or later.	Upgrade to the Cisco SAN-OS release required for the IVR features you want to use. See Table 8-1 and Chapter 2, "Troubleshooting Installs, Upgrades, and Reboots."

Host Does Not Have Write Access to Storage

Symptom Host does not have write access to storage.

Table 8-12 Host Does Not Have Write Access to Storage

Symptom	Possible Cause	Solution
Host does not have write access to storage.	Host is a member of a read-only zone.	If a host is a member of a read-only zone, the host has no write access to any IVR zone it may be a member of. Remove the host from the read-only zone.

Locked IVR CFS Session

IVR uses CFS to distribute the IVR configuration. If you enable IVR auto topology, it also uses CFS to distribute and update the IVR VSAN topology on all switches. In rare cases, you may encounter problems where CFS locks IVR so that you cannot modify the configuration.

Symptom Locked IVR CFS session.

Table 8-13 Locked IVR CFS Session

Symptom	Possible Cause	Solution
Locked IVR CFS session.	CFS did not give up the session lock for IVR after the last commit or an IVR configuration change is pending and has not been committed.	Choose **Fabric*xx* > All VSANS > IVR** and select the **CFS** tab. Set the ConfigView As drop-down menu to **p**ending and verify the pending configuration changes. Set the ConfigAction drop-down menu to **commit** to save these changes, **abort** to discard the changes, or **clear** to clear the session lock. Click **Apply Changes**. Or use the **show ivr pending-diff** CLI command to determine if you have a pending configuration change. Use **ivr commit** to commit this change or i**vr abort** to discard the changes and free up the session lock. If you do not have pending configuration changes, use the **clear ivr session** CLI command to free the session lock.

CFS Merge Failed

Symptom CFS merge failed.

Table 8-14 CFS Merge Failed

Symptom	Possible Cause	Solution
CFS merge failed.	IVR topology incorrect.	Choose **Fabricxx > All VSANs > IVR** and select the **Action** tab in Fabric Manager. Check the **Auto Discover Topology** check box and click **Apply Changes.** Select the **CFS** tab and set ConfigAction to **commit** and click **Apply Changes.** Or use either the **ivr vsan topology auto** CLI command to automatically reconfigure the IVR topology, or the **ivr vsan topology database** CLI command to manually reconfigure the IVR topology.
	Maximum number of VSANs or IVR VSAN topology entries reached.	Reconfigure your fabric before merging to reduce the number of VSANs or topology entries . See Table 8-1.
	Conflicting entries in the AF ID database.	Modify the conflicting entries in the AFID database.
	Conflicting user-configured IVR VSAN topology database entries.	Enable IVR auto topology on both fabrics before the merge and remove any user-configured IVR VSAN topology database entries.

Troubleshooting the IVR Wizard

The IVR Wizard in Fabric Manager simplifies the process of configuring IVR across your fabric. The IVR Wizard automatically checks for the appropriate Cisco SAN-OS version across the switches in the VSAN and determines which IVR features the switches are capable of. (See Table 8-1.)

This section describes the following warning or error dialog boxes that display when you configure IVR using the Fabric Manager IVR wizard:

Warning: Not All Switches Are IVR NAT Capable or Are Unmanageable

Symptom Warning: Not all switches are IVR NAT capable or are unmanageable.

Table 8-15 Not All Switches Are IVR NAT Capable or Are Unmanageable

Symptom	Possible Cause	Solution
Warning: Not all switches are IVR NAT capable or are unmanageable.	One or more switches in the fabric are not running Cisco MDS SAN-OS Release 2.1(1a) or later.	Upgrade to the Cisco SAN-OS release required for the IVR features you want to use. See Table 8-1 and Chapter 2, "Troubleshooting Installs, Upgrades, and Reboots."
	One or more switches in the fabric cannot communicate with Fabric Manager or are not Cisco SAN-OS switches.	Determine if any of the problem switches are required in the IVR topology. If not, ignore this message and proceed with the IVR configuration. If they are required, choose **Switches** and check the Status column to determine the cause and address the problem.

Error: The Following Switches Do Not Have Unique Domain IDs

Symptom The following switches do not have unique domain IDs.

Table 8-16 The Following Switches Do Not Have Unique Domain IDs

Symptom	Possible Cause	Solution
The following switches do not have unique domain IDs.	The listed switches have duplicate domain IDs in two or more VSANs in your proposed IVR configuration.	Choose **Fabric*xx* > All VSANS > Domain Manager** and set the ConfigDomainId to a unique number and set the Config Type drop-down menu to **static**. Set the Restart drop-down menu to **disruptive** and Click **Apply Changes**. This triggers a disruptive restart to make the running domain ID match the configured domain ID.
		Use IVR NAT. This may require upgrading to Cisco MDS SAN-OS Release 2.1(1a) or later.

Error: Pending Action/ Pending Commits

Symptom Pending action on pending commit error displays.

Table 8-17 Pending Action/Pending Commits

Symptom	Possible Cause	Solution
Pending action on pending commit error displays.	A separate IVR configuration change that was not committed.	IVR has pending changes that were not committed. Choose **Fabric*xx* > All VSANS > IVR** and select the **CFS** tab. Set the View Config As drop-down menu to **pending** and verify the pending configuration changes. Set the ConfigAction drop-down menu to **commit** to save these changes or **abort** to discard the changes**.** Click **Apply Changes**.
	The IVR CFS session was not unlocked after the last commit.	Choose **Fabric*xx* > All VSANS > IVR** and select the **CFS** tab. Set the ConfigAction drop-down menu to **clear** to remove the session lock. Click **Apply Changes**.

Error: Fabric Is Changing. Please Retry the Request Later

This error may occur where there are different versions of Cisco SAN-OS on the IVR-enabled switches. You should upgrade all IVR-enabled switches to the same version of Cisco SAN-OS.

CHAPTER 9

Troubleshooting Zones and Zone Sets

Zoning enables access control between storage devices and user groups. Creating zones increases network security and prevents data loss or corruption.

Zone sets consist of one or more zones. A zone set can be activated or deactivated as a single entity across all switches in the fabric, but only one zone set can be activated at any time.

Zones can be members of more than one zone set. A zone consists of multiple zone members. Members in a zone can access each other; members in different zones cannot access each other.

This chapter describes how to identify and resolve problems that might occur while implementing zones and zone sets on switches in the Cisco MDS 9000 Family. It includes the following sections:

Best Practices

This section provides the best practices for implementing zones and zone sets.

- Fibre Channel zoning should always be used.

  Creating zones increases network security and prevents data loss or corruption.

- Each host bus adapter (HBA) should have its own zone.

  In general, we recommend that the number of zones equal the number of HBAs communicating with the storage device. For example, if there are two hosts each with two HBAs communicating with three storage devices, we recommend using four zones. This type of zoning is sometimes referred to as *single initiator zoning*.

- Preplan your zone configuration, keeping in mind that multiple zone sets can be configured, but only one zone set can be active.
- Keep documented backups of zone members and zones within zone sets.
- Device aliases or FC aliases should be used to simplify management whenever possible.

  It is easier to identify devices with aliases than with WWNs. In general, you should assign aliases to WWNs.

- Use enhanced zoning whenever possible. Enhanced zoning is less disruptive, and ensures fabric-wide consistency for your zone configuration.

- Zone administration should generally be confined to a single Fibre Channel switch.

  Confining zone administration to a single Fibre Channel switch within a given fabric generally ensures that there is no possibility of activating an incomplete zone set, which could happen if the full zone set database is not consistent across Fibre Channel switches.

- The default zone policy should be deny (default).

  Leave the default zone policy as "deny" so that devices cannot inadvertently access each other when placed in the default zone.

- If using basic zoning, then choose **Fabric*xx* > VSAN*xx* > *zonesetname*** and select **FullZoneSet** from the Propagation drop-down menu in Fabric Manager. Or use the **zoneset distribute full vsan** CLI command to distribute the full zone database across the fabric whenever a zone set activation occurs. This ensures a consistent full zone database on all switches for that VSAN.

Troubleshooting Checklist

The following criteria must be met for zoning to function properly:

Checklist	Checkoff
Verify that you have an active zone set.	❑
Verify that you have the correct hosts and storage devices in the same zone.	❑
Verify that the zone is part of the active zone set.	❑
Verify that the default zone policy is permit if you are not using zoning.	❑

For zone configuration problems, use the following helpful tools:

- Cisco Fabric Analyzer. (Seethe "Cisco Fabric Analyzer" section on page B-23.)
- Cisco Fabric Manager and CLI system messages (See the System Messages, page 1-9.)
- Log messages (See the "Troubleshooting with Logs" section on page 1-12.)

Troubleshooting Zone Configuration Issues with Fabric Manager

Much of the information accessible through Fabric Manager can also be accessed using the CLI. (See the "Troubleshooting Zone Configuration Issues with the CLI" section on page 9-3.)

To verify which devices belong to the active zone set on a specific VSAN using Fabric Manager, follow these steps:

Step 1 Choose **Tools > Edit Full Zone Database** and select the **VSAN** from the drop-down menu. You see the full zone database for that VSAN. The active zone set appears in bold. If there is no zone set in bold, you have not activated a zone set for this VSAN.

Step 2 Expand the active zone set. You see the active zones displayed as new folders.

Step 3 Click on a zone. You see the devices belonging to the zone listed in the column on the left side of the dialog box. They are also highlighted in the map view.

Troubleshooting Zone Configuration Issues with the CLI

Much of the information accessed and summarized using the Fabric Manager can be found using CLI **show** commands. (See Table 9-1.)

Table 9-1 Zone Troubleshooting Commands in the CLI

Command	Command Description
show zone name *zonename*	Displays members of a specific zone.
show device-alias database	Displays any device aliases configured.
show fcalias *vsan-id*	Displays if and how FC aliases are configured.
show zone member *pWWN-id, fcalias-id, or pWWN-id*	Displays all zones to which a member belongs using the FC ID, the FC alias, or the pWWN.
show zone statistics	Displays the number of control frames exchanged with other switches.
show zone internal *vsan-id*	Displays the internal state of the zone server for a specific VSAN.
show zoneset *zonesetname*	Displays information about the named zone set.
show zoneset active	Displays information about the active zone set.

Note To issue commands with the **internal** keyword, you must have a network-admin group account.

The **debug zone change** CLI command followed by the zone name in question can help you get started debugging zones for protocol errors, events, and packets.

Note To enable debugging for zones, use the **debug zone** command in EXEC mode. To disable a debug command, use the **no** form of the command or use the **no debug all** command to turn off all debugging.

For protocol errors, use:

```
debug zone change errors vsan-id
```

For protocol events, use:

```
debug zone change events vsan-id
```

For protocol packets, use:

```
debug zone change packets vsan-id
```

Other useful **debug** commands include:

```
debug zone {all |
    change {errors | events | packets} |
    database {detail | errors | events} |
    gs errors {errors | events | packets} |
    lun-zoning {errors | events | packets} |
    merge {errors | events | packets} |
    mts notifications |
    pss {errors | events} ||
    read-only-zoning {errors | events | packets} |
    tcam errors {errors | events | packets} |
    transit {errors | events}} [vsan vsan-id]
```

Zone and Zone Set Issues

The section covers the following zone and zone set issues:

Host Cannot Communicate with Storage

A host cannot see a storage device for the following reasons:

- The default zone policy does not allow the devices to communicate.
- Storage devices and host interfaces do not belong to the same zone or the zone is not part of the active zone set.

Symptom Host cannot communicate with storage.

Table 9-2 Host Cannot Communicate with Storage

Symptom	Possible Cause	Solution
Host cannot communicate with storage.	Host and storage are not in the same zone.	See the "Resolving Host Not Communicating with Storage Issue Using Fabric Manager" section on page 9-4 or the "Resolving Host Not Communicating with Storage Using the CLI" section on page 9-6.
	Zone is not in active zone set.	
	No active zone set and default zone policy is deny.	
	The xE port connecting to the remote switch is isolated.	See the "xE Port Is Isolated in a VSAN" section on page 7-7.
	Host and storage are not in the same VSAN.	Verify the VSAN membership. See the "Verifying VSAN Membership Using Fabric Manager" section on page 7-6 or the "Verifying VSAN Membership Using the CLI" section on page 7-6.

Resolving Host Not Communicating with Storage Issue Using Fabric Manager

To verify that the host is not communicating with storage using Fabric Manager, follow these steps:

Step 1 Verify that the host and storage device are in the same VSAN. See the "Verifying VSAN Membership Using Fabric Manager" section on page 7-6.

Step 2 Configure zoning, if necessary, by choose **Fabric*xx* > VSAN*xx* > Default Zone** and selecting the **Policies** tab to determine if the default zone policy is set to **deny**.

The default zone policy of **permit** means all nodes can see all other nodes. **Deny** means all nodes are isolated when not explicitly placed in a zone.

Step 3 Optionally, select **permit** from the Default Zone Behavior drop-down menu to set the default zone policy to permit if you are not using zoning. Got to Step 8.

Step 4 Choose **Zone > Edit Local Full Zone Database** and select the VSAN you are interested in. Click on the zones folder and verify that the host and storage are both members of the same zone. If they are not in the same zone, see the "Resolving Host and Storage Not in the Same Zone Using Fabric Manager" section on page 9-5.

Step 5 Choose **Zone > Edit Local Full Zone Database** and select the VSAN you are interested in. Click on the active zone folder and determine if the zone in Step 5 and the host and disk appear in the active zone set. If the zone is not in the active zone set, see the "Resolving Zone is Not in Active Zone Set Using Fabric Manager" section on page 9-6.

Step 6 If there is no active zone set, right-click the zone set you want to activate in the Edit Local Full Zone Database dialog box and select **Activate** to activate the zone set.

Step 7 Verify that the host and storage can now communicate.

Resolving Host and Storage Not in the Same Zone Using Fabric Manager

To move the host and storage device into the same zone using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** and select the VSAN you are interested in. Click on the zones folder and find the zones that the host and storage are members of.

Step 2 Click on the zone that contains the host or storage that you want to move. Right-click on the row that represents this zone member and select **Delete** from the pop-up menu to remove this end device from the zone.

Step 3 Click on the zone that you want to move the end device to. Click and drag the row that represents the end device in the bottom table and add it to the zone in the top table.

Step 4 Verify that you have an active zone set for this VSAN by selecting the zone set name that appears in bold. If you do not have an active zone set, right-click on the zone set you want to activate in the Edit Local Full Zone Database dialog box and select **Activate** to activate the zone set.

Step 5 Expand the active zone set folder to verify that the zone in Step 3 is in the active zone set. If it is not, see the "Resolving Zone is Not in Active Zone Set Using Fabric Manager" section on page 9-6.

Step 6 Click **Activate...** to activate the modified zone set.

Step 7 Verify that the host and storage can now communicate.

Resolving Zone is Not in Active Zone Set Using Fabric Manager

To add a zone to the active zone set using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** and select the VSAN you are interested in. Right-click on the active zone set, which is in bold, and select **Insert**.

Step 2 Click on the zone that you want to add to this zone set and click **Add**.

Step 3 Click **Activate...** to activate the modified zone set.

Step 4 Verify that the host and storage can now communicate.

Resolving Host Not Communicating with Storage Using the CLI

To verify that the host is not communicating with storage using the CLI, follow these steps:

Step 1 Verify that the host and storage device are in the same VSAN. See the "Verifying VSAN Membership Using the CLI" section on page 7-6.

Step 2 Configure zoning, if necessary, by using the **show zone status** *vsan-id* command to determine if the default zone policy is set to **deny**.

```
switch# show zone status vsan 1
VSAN: 1 default-zone: deny distribute: active only Interop: default
    mode: basic merge-control: allow session: none
    hard-zoning: enabled
Default zone:
    qos: low broadcast: disabled ronly: disabled
Full Zoning Database :
    Zonesets:0  Zones:0 Aliases: 0
Active Zoning Database :
    Name: Database Not Available
Status:
```

The default zone policy of **permit** means all nodes can see all other nodes. **Deny** means all nodes are isolated when not explicitly placed in a zone.

Step 3 Optionally, use the **zone default-zone permit** CLI command to set the default zone policy to permit if you are not using zoning. Go to Step 7.

Step 4 Use the **show zone member** CLI command for host and storage device to verify that they are both in the same zone. If they are not in the same zone, see the "Resolving Host and Storage Not in the Same Zone Using Fabric Manager" section on page 9-5.

Step 5 Use the **show zoneset active** command to determine if the zone in Step 4 and the host and disk appear in the active zone set.

```
v_188# show zoneset active vsan 2
zoneset name ZoneSet3 vsan 2
  zone name Zone5 vsan 2
    pwwn 10:00:00:00:77:99:7a:1b [Hostalias]
    pwwn 21:21:21:21:21:21:21:21 [Diskalias]
```

Step 6 If the zone is not in the active zone set, see the "Resolving Zone is Not in Active Zone Set Using Fabric Manager" section on page 9-6.

Step 7 If there is no active zone set, use the **zoneset activate** command to activate the zone set.

```
switch(config)# zoneset activate ZoneSet1 vsan 2.
```

Step 8 Verify that the host and storage can now communicate.

Resolving Host and Storage Not in the Same Zone Using the CLI

To move the host and storage device into the same zone using the CLI, follow these steps:

Step 1 Use the **zone name** *zonename vsan-id* command to create a zone in the VSAN if necessary, and add the host or storage into this zone.

```
ca-9506(config)# zone name NewZoneName vsan 2
ca-9506(config-zone)# member pwwn 22:35:00:0c:85:e9:d2:c2
ca-9506(config-zone)# member pwwn 10:00:00:00:c9:32:8b:a8
```

Note The pWWNs for zone members can be obtained from the device or by issuing the **show flogi database** *vsan-id* command.

Step 2 Use the **show zone** command to verify that host and storage are now in the same zone.

```
switchA# show zone
zone name NewZoneName vsan 2
  pwwn 22:35:00:0c:85:e9:d2:c2
  pwwn 10:00:00:00:c9:32:8b:a8

zone name Zone2 vsan 4
  pwwn 10:00:00:e0:02:21:df:ef
  pwwn 20:00:00:e0:69:a1:b9:fc

zone name zone-cc vsan 5
  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
```

Step 3 Use the **show zoneset active** command to verify that you have an active zone set. If you do not have an active zone set, use the **zoneset activate** command to activate the zone set.

Step 4 Use the **show zoneset active** command to verify that the zone in Step 2 is in the active zone set. If it is not, use the **zoneset name** command to enter the zone set configuration submode, and use the **member** command to add the zone to the active zone set.

```
switch(config)# zoneset name zoneset1 vsan 2
ca-9506(config-zoneset)# member NewZoneName
```

Step 5 Use the **zoneset activate** command to activate the zone set.

```
switch(config)# zoneset activate ZoneSet1 vsan 2
```

Step 6 Verify that the host and storage can now communicate.

Resolving Zone is Not in Active Zone Set Using the CLI

To add a zone to the active zone set using the CLI, follow these steps:

Step 1 Use the **show zoneset active** command to verify that you have an active zone set. If you do not have an active zone set, use the **zoneset activate** command to activate the zone set.

Step 2 Use the **show zoneset active** command to verify that the zone in Step 1 is not in the active zone set.

Step 3 Use the **zoneset name** command to enter the zone set configuration submode, and use the **member** command to add the zone to the active zone set.

```
switch(config)# zoneset name zoneset1 vsan 2
ca-9506(config-zoneset)# member NewZoneName
```

Step 4 Use the **zoneset activate** command to activate the zone set.

```
switch(config)# zoneset activate ZoneSet1 vsan 2
```

Step 5 Verify that the host and storage can now communicate.

Troubleshooting Zone Set Activation

When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the *active zone set*. A zone that is part of an active zone set is called an *active zone*. Two main problems can occur with activating a zone set:

- No zone set is active.
- Zone set activation fails.

Zone activation can fail if a new switch joins the fabric. When a new switch joins the fabric, it acquires the existing zone sets. Also, large zone sets may experience timeout errors in Cisco MDS SAN-OS Release 1.3(4a) and earlier.

When a zone set activation fails, you may see the following system messages:

Error Message `ZONE-2-ZS_CHANGE_ACTIVATION_FAILED: Activation failed.`

Explanation The zone server cannot activate the zone set.

Recommended Action Use the **zoneset activate** CLI command or similar Fabric Manager procedure to.

Error Message `ZONE-2-ZS_CHANGE_ACTIVATION_FAILED_RESN: Activation failed : reason [chars].`

Explanation The zone server cannot activate because of reason shown in the error message.

Recommended Action No action is required.

Error Message `ZONE-2-ZS_CHANGE_ACTIVATION_FAILED_RESN_DOM: Activation failed : reason [chars] domain [dec].`

Explanation The zone server cannot activate because of reason shown in the error message on the domain.

Recommended Action No action is required.

Troubleshooting Zone Activation Using Fabric Manager

To verify the active zone set and active zones using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** and select the VSAN you are interested in. Click on the active zone set, which is in bold.

Step 2 Verify that the needed zones are active. If a zone is missing from the active zone set, see the "Resolving Zone is Not in Active Zone Set Using Fabric Manager" section on page 9-6.

Step 3 Click **Activate...** to activate the zone set.

Step 4 If you are still experiencing zone set activation failure, use the **show zone internal change event-history vsan <*vsan-id*>** CLI command to determine the source of zone set activation problem.

Troubleshooting Zone Activation Using the CLI

To verify the active zone set and active zones using the CLI, follow these steps:

Step 1 Use the **show zoneset active** *vsan-id* command to display the active zones.

```
switchA# show zoneset active vsan 2
zoneset name ZoneSet1 vsan 2
  zone name NewZoneName vsan 2
    * pwwn 22:35:00:0c:85:e9:d2:c2
    * pwwn 10:00:00:00:c9:32:8b:a8
```

Step 2 Verify that the needed zones are active.

Step 3 Optionally, use the **zoneset name** *ActiveZonesetName vsan-id* command and the **member** *NewZone* command to add the zone to the active zone set in the VSAN.

```
switch(config)# zoneset name ZoneSet1 vsan 2
switch(config-zoneset)# member NewZoneAdded
```

Step 4 Use the **zoneset activate** command to activate the zone set.

```
switch(config)# zoneset activate ZoneSet1 vsan 2
```

Step 5 If you are still experiencing zone set activation failure, use the **show zone internal change event-history vsan <*vsan-id*>** command to determine the source of the zone set activation problem.

Troubleshooting Full Zone Database Synchronization Across Switches

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

Resolving Out of Sync Full Zone Database Using Fabric Manager

To verify if the full zone database is in sync across switches using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > VSAN*xx* > *zonesetname*** and select the **Policies** tab.

Step 2 Verify that the Propagation field is set to **FullZoneSet**. If it is not, select **FullZoneSet** from the drop-down menu.

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

Resolving an Out of Sync Full Zone Database Using the CLI

To verify if the full zone database is in sync across switches using the CLI, follow these steps:

Step 1 Use the **show zone status** command to verify if the distribute flag is on.

```
switch# config t show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
    mode: basic merge-control: allow session: none
    hard-zoning: enabled
Default zone:
    qos: low broadcast: disabled ronly: disabled
Full Zoning Database :
    Zonesets:3  Zones:7 Aliases: 9
Active Zoning Database :
    Name: ZoneSet1  Zonesets:1  Zones:2
Status:
```

This example shows that only the active zone set is distributed.

Step 2 Verify that the distribute flag is on.

Mismatched Default Zone Policy

If you are using basic zoning, you must verify that the default zone policy is the same for all switches in the VSAN. If the default zone policy varies, then you may experience zoning problems. If all switches in the VSAN have Cisco SAN-OS Release 2.0(1b) or later, you can use enhanced zoning. Enhanced zoning synchronizes your zone configuration across all switches in the VSAN, eliminating the possibility of mismatched default zone policies.

Resolving Mismatched Default Zone Policies Using Fabric Manager

To resolve mismatched default zone policies using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > VSAN*xx* > *zonesetname*** and select the **Policies** tab.

Step 2 View the Default Zone Behavior field for each switch in the VSAN to determine which switches have mismatched default zone policies.

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

Step 4 If you are using basic zoning, Select the same value from the Default Zone Behavior drop-down menu for each switch in the VSAN to set the same default zone policy.

Step 5 If you are using enhanced zoning, follow these steps:

a. Choose **Fabric*xx* > VSAN*xx*** and view the Release field to verify that all switches are capable of working in the enhanced mode.
All switches must have Cisco MDS SAN-OS Release 2.0(1b) or later. If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.

b. Choose **Fabric*xx* > VSAN*xx* > *zonesetname*** and select the **Policies** tab and set Default Zone Behavior field to set the default zone policy.

c. Click **Apply Changes** to save these changes.

d. Select the **Enhanced** tab and select **enhanced** from the Action drop-down menu.

e. Click **Apply Changes** to save these changes.
By doing so, you automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies, and then release the lock. All switches in the VSAN then move to the enhanced zoning mode.

Note After moving from basic zoning to enhanced zoning (or vice versa), we recommend that you save the running configuration.

Resolving Mismatched Default Zone Policies Using the CLI

To resolve mismatched default zone policies using the CLI, follow these steps:

Step 1 Issue the **show zone status** command.

```
v_188# show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
    mode: basic merge-control: allow session: none  <------------------
    hard-zoning: enabled
Default zone:
    qos: low broadcast: disabled ronly: disabled
Full Zoning Database :
    Zonesets:5  Zones:18 Aliases: 11
Active Zoning Database :
    Name: ZoneSet1  Zonesets:1  Zones:2
Status:
```

This example shows the default zone policy is deny, and the zone mode is basic.

Step 2 If you are using basic zoning, follow these steps:

a. Repeat Step 1 for all switches in the VSAN to verify that they have the same zone mode. Use the **zone mode basic** command to change any switches that are not in basic mode.

b. Use the **zone default-zone** command on each switch in the VSAN to set the same default zone policy.

Step 3 If you are using enhanced zoning, follow these steps:

a. Use the **show version** command on all switches in the VSAN to verify that all switches are capable of working in the enhanced mode.
All switches must have Cisco MDS SAN-OS Release 2.0(1b) or later. If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.

b. Use the **zone default-zone** command to set the default zone policy.

c. Use the **zone mode enhanced** *vsan-id* command to set the operation mode to enhanced zoning mode.
By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies, and then release the lock. All switches in the VSAN then move to the enhanced zoning mode.

```
switch(config)# zone mode enhanced vsan 3000
```

Note After moving from basic zoning to enhanced zoning (or vice versa), we recommend that you use the **copy running-config startup-config** command to save the running configuration.

Zone Merge Failure

A zone merge request may fail because of the following configuration issues:

- Too many zone sets
- Too many aliases

- Too many attribute groups
- Too many zones
- Too many LUN members
- Too many zone members

Use the **show zone internal merge event-history** CLI command to determine the cause of the zone merge failure.

You may see one or more of the following system messages after a zone merge failure:

Error Message `ZONE-2-ZS_MERGE_ADJ_NO_RESPONSE: Adjacent switch not responding, Isolating Interface [chars] (VSAN [dec]).`

Explanation Interface on the VSAN was isolated because the adjacent switch is not responding to zone server requests.

Recommended Action Flap the interface.

Introduced Cisco MDS SAN-OS Release 1.2(2a).

Error Message `ZONE-2-ZS_MERGE_FAILED: Zone merge failure, Isolating interface [chars].`

Explanation Interface isolated because of a zone merge failure.

Recommended Action Compare active zoneset with the adjacent switch or enter the **zone merge interface** CLI command or similar Fabric Manager/Device Manager command.

Introduced Cisco MDS SAN-OS Release 1.2(2a).

Error Message `ZONE-2-ZS_MERGE_FULL_DATABASE_MISMATCH: Zone merge full database mismatch on interface [chars].`

Explanation Full zoning databases are inconsistent between two switches connected by interface . Databases are not merged.

Recommended Action Compare full zoning database with the adjacent switch. Correct the difference and flap the link.

Introduced Cisco MDS SAN-OS Release 1.3(1).

Error Message `ZONE-2-ZS_MERGE_FULL_DATABASE_MISMATCH: Zone merge full database mismatch on interface [chars].`

Explanation Full zoning databases are inconsistent between two switches connected by the interface. Databases are not merged.

Recommended Action Compare full zoning database with the adjacent switch, correct the difference and flap the link.

Introduced Cisco MDS SAN-OS Release 1.2(2a).

Error Message `ZONE-2-ZS_MERGE_UNKNOWN_FORMAT: Unknown format, isolating interface [chars].`

Explanation Interface isolated because of an unknown format in the merge request.

Recommended Action Set the interoperability mode to the same value on both switches.

Introduced Cisco MDS SAN-OS Release 2.0(1b).

Note Zoning information exists on a per VSAN basis. Therefore, for a TE port, it may be necessary to verify that the zoning information does not conflict with any allowed VSAN.

Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, the port may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set.
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

If after verifying the Fibre Channel name server , you still experience FSPF problems (such as discovering remote switches and their attached resources), the fabric may have zone configuration problems. Examples of zone configuration problems are mismatched active zone sets and misconfigured zones within the active zone set.

Resolving a Link Isolation Because of a Failed Zone Merge Using Fabric Manager

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

To perform a zone merge analysis using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Merge Analysis** from the Zone menu.

You see the Zone Merge Analysis dialog box.

Step 2 Select the first switch to be analyzed from the Check Switch 1 drop-down list.

Step 3 Select the second switch to be analyzed from the And Switch 2 drop-down list.

Step 4 Enter the VSAN ID where the zone set merge failure occurred in the For Active Zoneset Merge Problems in VSAN Id field.

Step 5 Click **Analyze** to analyze the zone merge. Click **Clear** to clear the analysis data from the Zone Merge Analysis dialog box.

Resolving a Link Isolation Because of a Failed Zone Merge Using the CLI

The following CLI commands are used to resolve a failed zone merge:

- **zoneset import** *vsan-id*
- **zoneset export** *vsan-id*

To resolve a link isolation because of a failed zone merge using the CLI, follow these steps:

Step 1 Use the **show interface** command to confirm that the port is isolated because of a zone merge failure.

```
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
```

An E port is segmented (isolation due to zone merge failure) if the following conditions are true:

- The active zone sets on the two switches differ from each other in terms of zone membership (provided there are zones at either side with identical names).
- The active zone set on both switches contain a zone with the same name but with different zone members.

Step 2 Verify the zoning information, using the following commands on each switch:

- **show zone vsan** *vsan-id*
- **show zoneset vsan** *vsan-id*

Step 3 You can use two different approaches to resolve a zone merge failure by overwriting the zoning configuration of one switch with the other switch's configuration. This can be done with either of the following commands:

- **zoneset import interface** *interface-number* **vsan** *vsan-id*
- **zoneset export interface** *interface-number* **vsan** *vsan-id*

The **import** option of the command overwrites the local switch's active zone set with that of the remote switch. The **export** option overwrites the remote switch's active zone set with the local switch's active zone set.

Step 4 If the zoning databases between the two switches are overwritten, you cannot use the **import** option. To work around this, you can manually change the content of the zone database on either of the switches, and then issue a **shutdown/no shutdown** command sequence on the isolated port.

Step 5 If the isolation is specific to one VSAN and not on an 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 reinsert it.

Note Do not simply issue a **shutdown/no shutdown** command sequence on the port. This would affect all the VSANs crossing the EISL instead of just the VSAN experiencing the isolation problem.

Mismatched Active Zone Sets Within the Same VSAN

When merging switch fabrics, you must ensure that the zones in both active zone sets 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.

For example, two switches may have the same zone set name, and the same zone names, but different zone members. As a result, the VSAN is isolated on the TE port that connects the two switches.

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

- Modify the zone members on both zone sets to match and eliminate the conflict.
- Deactivate the zone set on one of the switches and restart the zone merge process.
- Explicitly import or export a zone set between the switches to synchronize them.

Resolving Mismatched Active Zone Sets Within the Same VSAN Using Fabric Manager

Mismatched active zone sets within the same VSAN result in that VSAN being segmented in Fabric Manager. To verify a mismatched active zone set within the same VSAN using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Edit Local Full Zone Database** and select the segmented VSAN you are interested in. Click on the active zone set, which is in bold, to view the list of zones and zone members for this active zone set.

Step 2 Repeat Step 1 for the other segmented VSAN.

A mismatched active zone set may include zones with the same name but different members, or a missing zone within the zone set.

Step 3 Do one of the following to resolve the isolation problem:

- Change the membership of one of the zones to match the other zone of the same name. See the "Resolving Host and Storage Not in the Same Zone Using Fabric Manager" section on page 9-5.
- Discard one of the zone sets completely by deactivating it using the **no zoneset activate** command. If a VSAN does not have an active zone set, it automatically takes the active zone set of the other merging switch. See the "Deactivating a Zone Set and Restarting the Zone Merge Process Using Fabric Manager" section on page 9-19.
- Choose **Zone > Copy Full Zone Database** to overwrite the active zone set on one switch. This method is destructive to one of the active zone sets.

Resolving Mismatched Active Zone Sets Within the Same VSAN Using the CLI

To verify a mismatched active zone set within the same VSAN using the CLI, follow these steps:

Step 1 Use the **show zoneset active** *vsan-id* command to display the active zone set configuration of the first switch.

```
Switch1# show zoneset active vsan 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 Use the **show zoneset active** *vsan-id* command to display the active zone set configuration of the second switch:

```
Switch2# show zoneset active vsan 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 Issue the **show interface** command to view information about the TE port and the interface.

```
Switch2# show interface 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 Use the **show port internal interface** *interface number* CLI command to get information about why the interface is isolated.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

```
switch# show port internal info interface fc1/8

fc1/8 - if_index: 0x0109C000, phy_port_index: 0x3c
  Admin Config - state(up), mode(TE), 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(0x3)
    description()
    Hw Capabilities: 0xb
    trunk vsans (up) (1)
    .
    .
    .
    trunk vsans (isolated) (99)
  TE port per vsan information
  fc2/29, Vsan 1 - state(up), state reason(None), fcid(0x690202)
    port init flag(0x38000), current state [TE_FSM_ST_E_PORT_UP]
  fc2/29, Vsan 99 - state(down), state reason(Isolation due to zone merge failure),
fcid(0x000000)
        port init flag(0x0), current state [TE_FSM_ST_ISOLATED_VSAN_MISMATCH]
```

From this output, you can see the VSAN is isolated because of o a zone merge failure.

Step 5 Do one of the following to resolve the isolation problem:

- Change the membership of one of the zones to match the other zone of the same name. See the "Resolving Host and Storage Not in the Same Zone Using Fabric Manager" section on page 9-5.
- Discard one of the zone sets completely by deactivating it using the **no zoneset activate** command. If a VSAN does not have an active zone set, it automatically takes the active zone set of the other merging switch. See the "Deactivating a Zone Set and Restarting the Zone Merge Process Using the CLI" section on page 9-20.
- Overwrite the active zone set on one switch using the **import** or **export** commands. This method is destructive to one of the active zone sets.
  - **zoneset import interface** *interface-number* **vsan** *vsan-id*
  - **zoneset export interface** *interface-number* **vsan** *vsan-id*

Step 6 Use the **show interface** *fcx/y* **trunk** *vsan-id* command to verify that VSAN 99 is no longer isolated:

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

Deactivating a Zone Set and Restarting the Zone Merge Process Using Fabric Manager

To deactivate a zone set and restart the zone merge process using Fabric Manager, follow these steps:

Step 1 Choose **Zone > Deactivate** Zone Set to deactivate the zone set configuration.

Caution This will disrupt traffic and cause the MDS 9000 switch to lose connectivity with the network.

Step 2 Choose **Interfaces > FC Physica**l and select **down** from the Status Admin drop-down menu to shut down the connection to the zone to be merged. You may see the following system messages:

```
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/14 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/15 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/16 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 fc1/16 to none
```

Step 3 Choose **Interfaces > FC Physica**l and select **up** from the Status Admin drop-down menu to enable the connection to the zone to be merged. You may see the following system messages:

```
Nov 19 10:28:11 switch4 %LOG_PORT_CHANNEL-5-FOP_CHAN
GED: port-channel 1: first operational port changed from none to fc1/15
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 fc1/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/16, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/16, vsan 1 is up
```

Step 4 Choose **Zone > Edit Local Full Zone Database** to verify the active zone set configuration.

After deactivating the zone set onthe first switch and performing a shutdown followed by a no shutdown on the ISL that connects it to the second switch, the zone merge is processed again. Because the first switch has no active zone set, it learns the active zone set from the second switch during the zone merge process.

Deactivating a Zone Set and Restarting the Zone Merge Process Using the CLI

To deactivate a zone set and restart the zone merge process using the CLI, follow these steps:

Step 1 Use the **no zoneset activate name** *zoneset-name vsan-id* command to deactivate the zone set configuration from the switch:

Caution This will disrupt traffic and cause the MDS 9000 switch to lose connectivity with the network.

```
switch4(config)# no zoneset activate name excal2 vsan 1
Zoneset Deactivation initiated. check zone status
```

Step 2 Use the **show zoneset active** command to confirm that the zone set has been removed.

Step 3 Use the **shut down** command to shut down the connection to the zone to be merged.

```
switch4(config)# interface port-channel 1
switch4(config-if)# shutdown
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/14 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/15 is down
(Channel admin down)
Nov 19 10:26:10 switch4 %LOG_PORT-5-IF_DOWN_CHANNEL_ADMIN_DOWN: Interface fc1/16 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 fc1/16 to none
```

Step 4 Use the **no shutdown** command to reactivate the connection to the zone to be merged:

```
switch4(config-if)# no shutdown
Nov 19 10:28:11 switch4 %LOG_PORT_CHANNEL-5-FOP_CHAN
GED: port-channel 1: first operational port changed from none to fc1/15
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 fc1/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/16, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/14, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/15, vsan 1 is up
Nov 19 10:28:21 switch4 %LOG_PORT-5-IF_TRUNK_UP: Interface fc1/16, vsan 1 is up
```

Step 5 Use the **show zoneset** active *vsan-id* commands to exit configuration mode and check the active zone sets.

```
switch4# show zoneset active
zoneset name wall vsan 1
  zone name excal1 vsan 1
  * fcid 0x620200
    fcid 0x6200ca
zone name $default_zone$ vsan 1
  * fcid 0x6e00da
  * fcid 0x6e00d9
  * fcid 0x6e00d6
  * fcid 0x6e0100
```

After deactivating the zone set on switch 4 and performing a shutdown followed by a no shutdown on the ISL that connects it to switch 3, the zone merge is processed again. Because switch 3 has no active zone set, it learns the active zone set from switch 4 during the zone merge process.

Enhanced Zoning Issues

Enhanced zoning uses a session locking facility like CFS to prevent simultaneous zoning configuration changes by two users on the same or separate switches. When a user starts to make a zoning change on one switch for a VSAN, that switch will lock the fabric to prevent others from making zoning changes. The user must issue a commit to make the changes active and release the fabric wide lock.

Problems can occur when the lock is acquired, but not released. In this situation, you cannot configure zoning on that VSAN. If you are using the CLI, you see error messages when you attempt to enter the zoning configuration mode.

Troubleshooting CLI commands to use for enhanced zoning issues:

- **show zone internal change event-history**
- **show zone status vsan**
- **show zone pending-diff**
- **show zone pending vsan**

Symptom Cannot configure zoning.

Table 9-3 Cannot Configure Zoning

Symptom	Possible Causes	Solutions
Cannot configure zoning.	Another user on the same switch is holding the enhanced zoning configuration lock. If you are using the CLI, you see a message stating that another session is active.	See the"Resolving Enhanced Zoning Lock Issues with Fabric Manager" section on page 9-22 or the "Resolving Enhanced Zoning Lock Issues with the CLI" section on page 9-22.
	Another user on a different switch is holding the enhanced zoning configuration lock. If you are using the CLI, you see a message stating that the lock is currently busy.	

Resolving Enhanced Zoning Lock Issues with Fabric Manager

To resolve a lock failure using Fabric Manager, follow these steps:

Step 1 Choose **Fabric*xx* > VSAN*xx*** and select the zone set that you want to configure.

Step 2 Select the **Enhanced** tab from the Information pane and view the Config DB Locked By column to determine which switch and which user holds the enhanced zoning lock for this VSAN.

Step 3 Check the **Config DB Discard Changes** check box and click **Apply Changes** to clear the enhanced zoning lock.

Note Verify that no valid configuration change is in progress before you clear a lock

Resolving Enhanced Zoning Lock Issues with the CLI

To resolve a lock issue using the CLI, follow these steps:

Step 1 Use the **show zone internal vsan** CLI command to determine which switch has the lock for the VSAN.

```
switch# show zone internal vsan 16
VSAN: 16 default-zone: deny(rw) distribute: active only
    E_D_TOV: 2000 R_A_TOV: 10000 D_S_TOV: 5000 F_S_TOV: 5000 F_D_TOV: 2000
    Interop: default IOD: disable bcast: enable dflt-bcast: disable dflt-qos: 0
    DBLock:-(F count:0) Ifindex Table Size: 2 Transit Frame Index: 0
    Total Transit Frame Count: 11 Transit Discard Count: 9 Global Full Database Counters :
    Zonesets: 9 Zones: 153
    Aliases: 58 Attribute-groups: 15
    Members: 482 LUN Members: 0
Global Active Database Counters :
    Zones: 159 Members: 442 LUN Members: 0 Global Database (Active + Full) Counters :
     Read-only Zones: 0 LUN Members: 0
License Info: 0x0
Full Zoning Database :
    Zonesets:2  Zones:2 Aliases: 0 Attribute-groups: 1 Active Zoning Database :
    Name: CX400-BLUE  Zonesets:1  Zones:2 TCAM Info :
    cur_seq_num : 2840,  state : 0
    add_reqs = 15, del_reqs = 0, entries_added = 9 Change protocol info :
    local domain id = 50,   ACA by 0x58   <===========domain ID 58 has the lock
    State =       Idle,  reply_cnt = 1, req_sent_cnt = 1, req_pending =0
    Remote domains :
        58
```

Note If you see ACA by 0xff in the display, it means that no lock is known to exist on the domain in this switch. This should be the same for all switches in the VSAN.

Step 2 Use the **show zone status vsan** CLI command on the switch that holds the lock to determine the lock holder In the example above, you use this command on the switch that has the domain ID 58.

```
switch#show zone status vsan 16

  VSAN: 16 default-zone: deny distribute: active only Interop: default
    mode: enhanced merge-control: allow session: cli [admin]  <---- user admin has lock
```

```
hard-zoning: enabled
```

Step 3 Use the **no zone commit vsan** CLI command to release the lock if you are the holder of the lock.

Step 4 Use the **no zone commit vsan** <*vsan id*> **force** CLI command to release the lock if another user holds the lock.

Note Verify that no valid configuration change is in progress before you clear a lock

CHAPTER 10

Troubleshooting IP Storage Services

This chapter describes how to identify and resolve problems that might occur in the IP storage services portion of the Cisco MDS 9000 Family products. It includes the following sections:

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 10-1). 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 10-1 Connecting MDS 9000 Family Switches Over IP

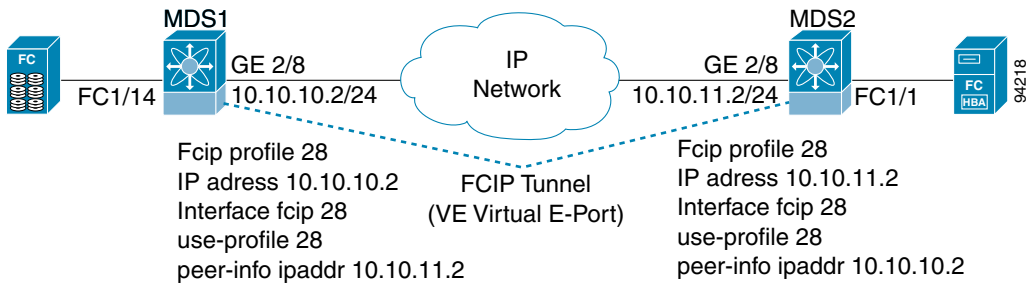

IP Connections Troubleshooting

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 802.1q, EtherChannel, and VRRP for iSCSI.

Verifying Basic Connectivity with the CLI

To verify basic connectivity using the CLI, follow these steps:

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 **show ip route**, **traceroute**, and **show arp** commands. A sample output of the **show ip route** command follows:

```
switch # show 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 interface GigE, verified using the **show 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 interface mgmt0, verified using the **show 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
```

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

```
switch# show ips arp interface gigabitethernet 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 interface gigabitethernet 4/7
arp clear successful
```

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

```
switch# show ips arp interface gigabitethernet 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 **show ips arp** command follows.

```
switch# show ips arp interface 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 **show arp** command follows.

```
switch# show 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
      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
```

Verification of Switch Connectivity

You can verify connectivity to a destination switch.

Note The FC ID variable used in this procedure is the domain controller address; it is not a duplication of the domain ID.

Verifying Switch Connectivity with the CLI

To verify connectivity to a destination switch, follow these steps:

	Command	Purpose

Step 1	`switch# show fcdomain domain-list vsan 200`
	`Number of domains: 7`
	`Domain ID WWN`
	`--------- -----------------------`
	`0x01(1) 20:c8:00:05:30:00:59:df [Principal]`
	`0x02(2) 20:c8:00:0b:5f:d5:9f:c1`
	`0x6f(111) 20:c8:00:05:30:00:60:df`
	`0xda(218) 20:c8:00:05:30:00:87:9f [Local]`
	`0x06(6) 20:c8:00:0b:46:79:f2:41`
	`0x04(4) 20:c8:00:05:30:00:86:5f`
	`0x6a(106) 20:c8:00:05:30:00:f8:e3`
		Displays the destination switch's domain ID.
		To obtain the domain controller address, concatenate the domain ID with FFFC. For example, if the domain ID is 0xda(218), the concatenated ID is 0xfffcda.

Step 2	`switch# fcping fcid 0xFFFCDA vsan 200`
	`28 bytes from 0xFFFCDA time = 298 usec`
	`28 bytes from 0xFFFCDA time = 260 usec`
	`28 bytes from 0xFFFCDA time = 298 usec`
	`28 bytes from 0xFFFCDA time = 294 usec`
	`28 bytes from 0xFFFCDA time = 292 usec`
	
	`5 frames sent, 5 frames received, 0 timeouts`
	`Round-trip min/avg/max = 260/288/298 usec`
		Verifies reachability of the destination switch by checking its end-to-end connectivity.

Verification of Static IP Routing

Static routing is a mechanism to configure IP routes on the switch.

Verifying Static IP Routing with the CLI

To verify that the IP routes are still there, use the **show ip route** command. A sample output of the **show ip route** command follows.

```
switch# show 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
```

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

```
switch# show 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
```

FCIP Connections Troubleshooting

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 10-2 shows the one-to-one topology used for configuration.

Figure 10-2 One-to-One Topology

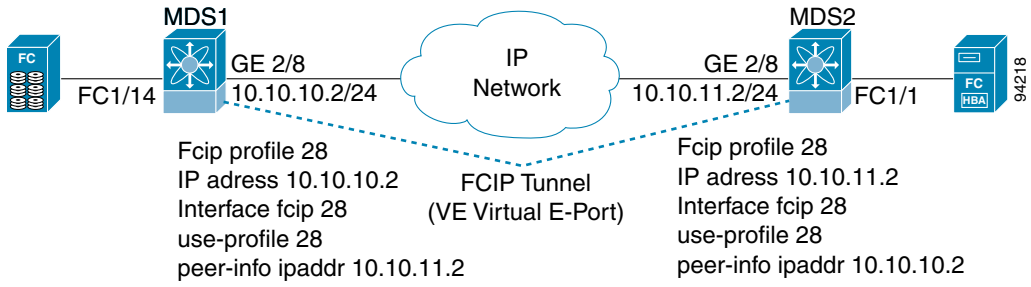

First, perform the following steps to configure the MDS1.

Configuration the First Switch with the CLI

To configure the first switch using the CLI, follow these steps:

Step 1 Enter configuration mode

Step 2 Set the interface

```
MDS1(config)# interface gigabitethernet 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
```

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 shutdown
MDS1(config-if)# end
```

Displaying the Default Values with the CLI

The output from issuing the **show running-config** command displays the default values in the following example.

```
MDS1# show running-config

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
```

Setting the Static Route for FCIP Tunnels with the CLI

The static route must be set for FCIP tunnels. This route could also be **ip route 10.10.11.0 255.255.255.0 interface gigabitethernet 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
```

Debugging the Configuration of the Second Switch with the CLI

The following example shows the configuration of a switch, MDS2, with debug mode activated. To activate debug mode for this situation, run the **debug ips flow fcip** command on a separate terminal.

```
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

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
```

```
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 shutdown
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)
```

Displaying the Debug Output from FCIP Tunnel Supervisor with the CLI

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

```
MDS2(config)# interface fcip 28
MDS2(config-if)# no shutdown
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
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
```

Displaying the Debug Output from the FCIP Tunnel IPS Module with the CLI

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 callbeck, 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...
```

Verifying the Configuration of the Profiles with the CLI

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
```

Verifying the Establishment of the FCIP Tunnel with the CLI

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.)

```
      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
                         Mode
-------------------------------------------------------------------------------
fcip28     1      auto   on      trunking     TE     28      --

MDS1# show interface fcip 28 counters brief

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ------------------------------ ------------------------------
                   Rate     Total                 Rate     Total
                   Mbits/s  Frames                Mbits/s  Frames
-------------------------------------------------------------------------------
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.)

Verifying the Establishment of Default TCP Connections for Each Configured FCIP Tunnel with the CLI

Verify two default 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

    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
```

Verifying the Statistics of the ASIC Chip on Each Gigabit Ethernet Port with with the CLI

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]
```

Ethereal Screen Captures of the TCP Connection and FCIP Tunnels

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).

Figure 10-3 First Capture of TCP Connection

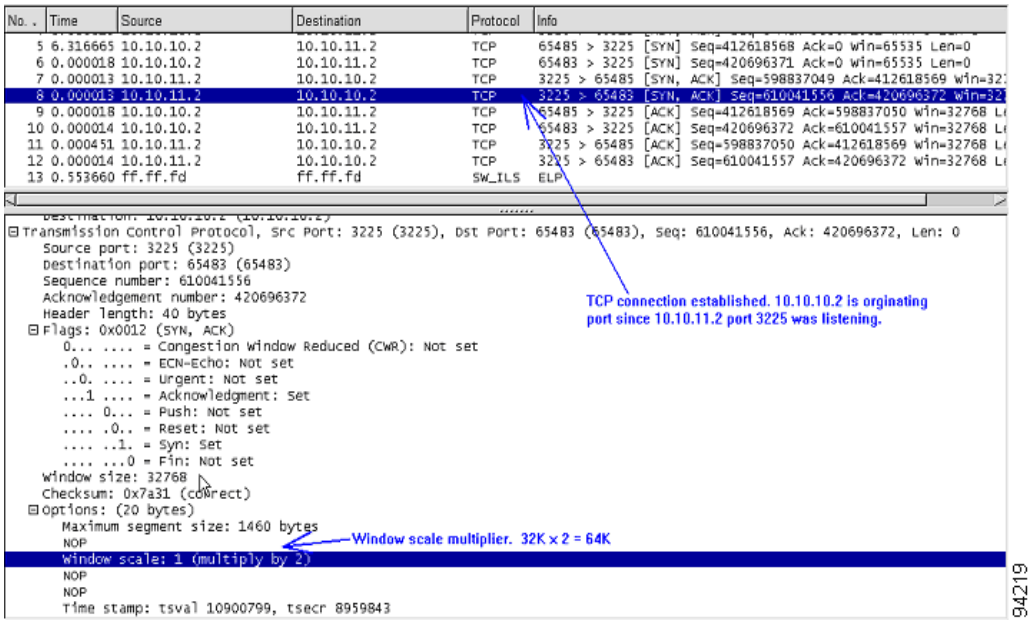

Figure 10-4 shows more of the trace, with frame 13 being the first FCIP frame. This frame carries the FC Standard ELP.

Figure 10-4 Second Capture of TCP Connection

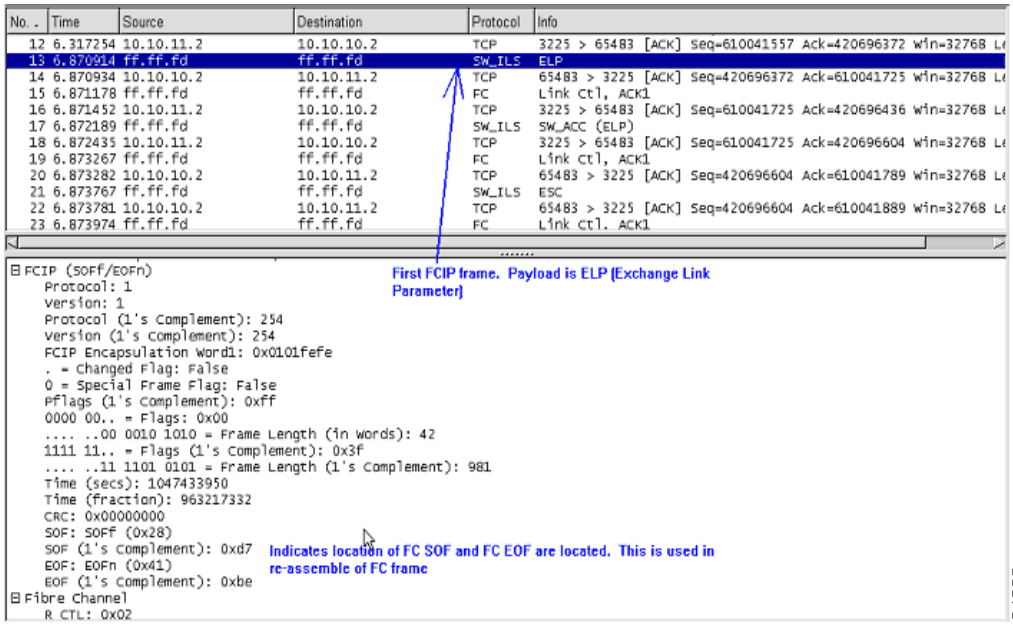

Figure 10-5 shows the FC portion of the EISL initialization over the FCIP tunnel.

Figure 10-5 Third Capture of TCP Connection

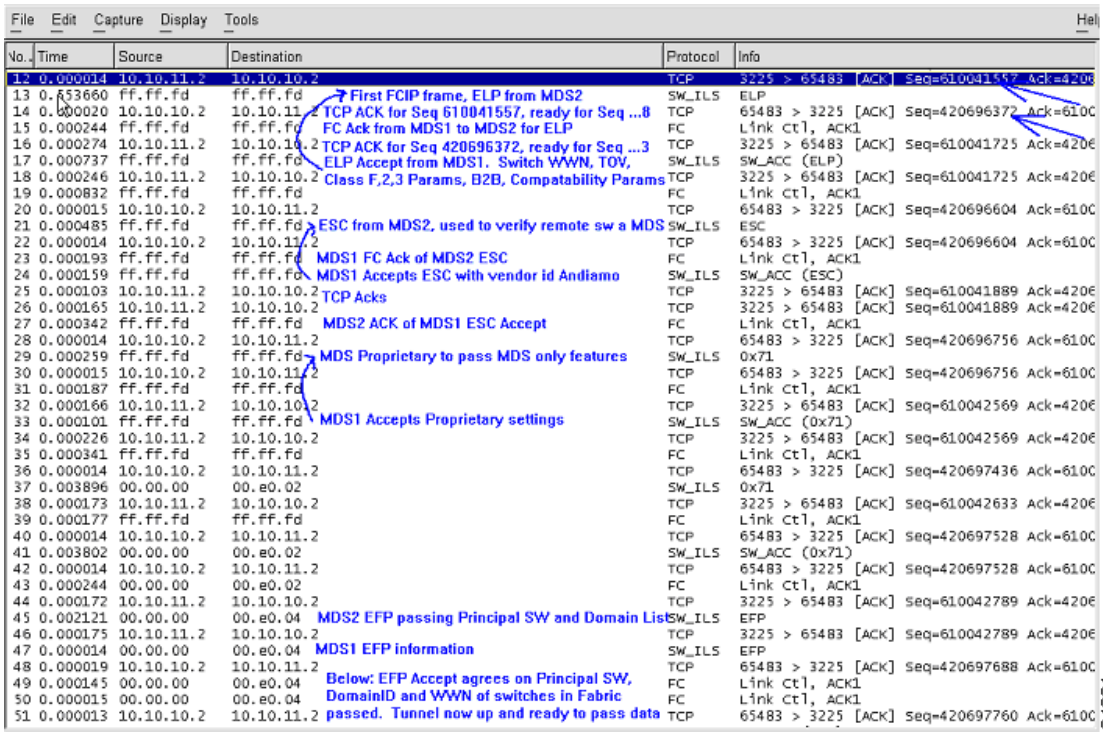

One-to-Three FCIP Tunnel Creation and Monitoring

Figure 10-6 shows the configuration of switch MDS1 for three tunnels from one Gigabit Ethernet port.

Figure 10-6 MDS1 Configured for Three FCIP Tunnels

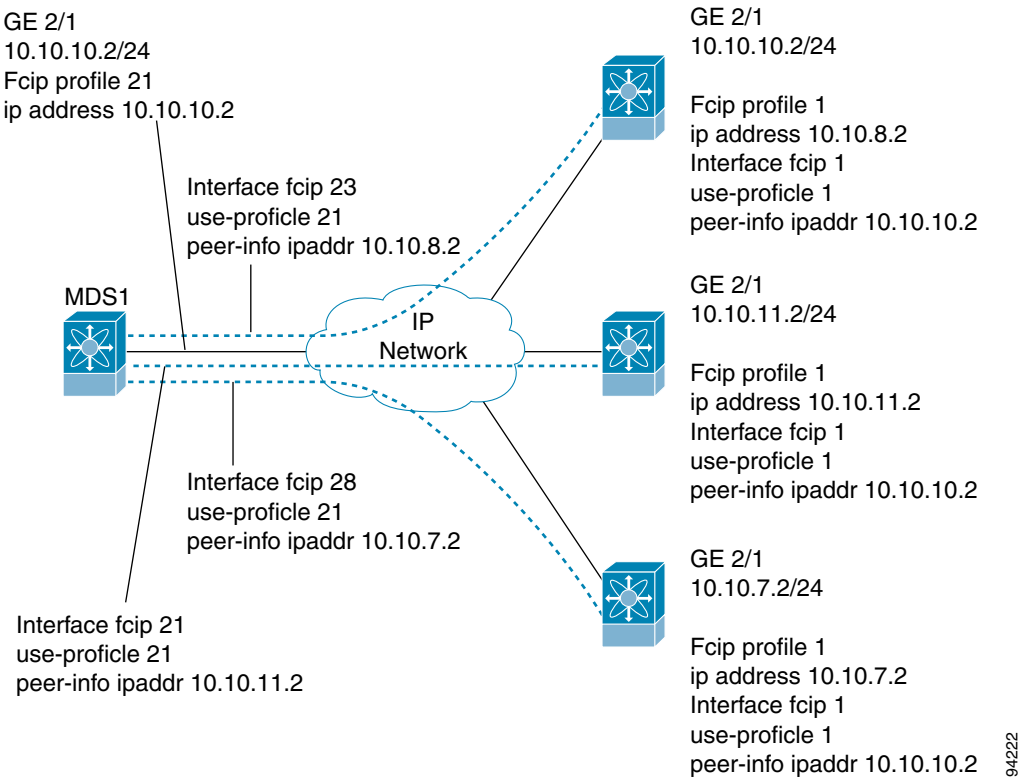

Displaying the Configuration of the First Switch with the CLI

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 shutdown
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 interface gigabitethernet 2/1
```

Creating the FCIP Interface for the Second Tunnel with the CLI

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)# interface fcip 23
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.8.2
```

```
MDS1(config-if)# no shutdown
MDS1(config-if)# exit
MDS1(config)#
```

Now the FCIP interface is created for the third tunnel.

```
MDS1(config)# interface 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
MDS1(config)#
```

FCIP Profile Misconfiguration Examples

The following example shows an incorrect or non existent IP address used for an FCIP profile.

Displaying Incorrect or Non-existent IP Address for Use with FCIP Profile with the CLI

```
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# config t
Enter configuration commands, one per line.  End with CNTL/Z.
MDS22(config)# interface gigabitethernet 2/5
MDS22(config-if)# ip addr 34.34.34.34 255.255.255.0
MDS22(config-if)# no shutdown
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 KBThe 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 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
```

Displaying Configuration Errors when Bringing Up a Tunnel on a Selected Port with the CLI

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)# interface fcip 21
MDS2(config-if)# passive-mode
```

(Put interface FCIP 21 in passive mode to guarantee MDS1 initiates a TCP connection.)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)# interface 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 interface 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

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.

Displaying FCIP Misconfiguration Examples with the CLI

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 interface 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
```

Displaying the Interface FCIP Shut Down Administratively with the CLI

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 interface 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 interface 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

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
```

Displaying the Debug Output from the Second Switch with the CLI

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 interface 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 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
```

Displaying Passive Mode Set on Both Sides of the FCIP Tunnel with the CLI

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 interface 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 interface 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)# interface fcip 21
MDS2(config-if)# no passive-mode
```

(Change one or both sides to **no passive-mode**.)

```
MDS2# show interface fcip 21
fcip21 is trunking
```

Displaying a Time Stamp Acceptable Difference Failure with the CLI

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)# interface 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.)

```
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 interface 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
```

Figure 10-7 shows a trace of timestamp difference failure.

Figure 10-7 Trace of Timestamp Difference Failure

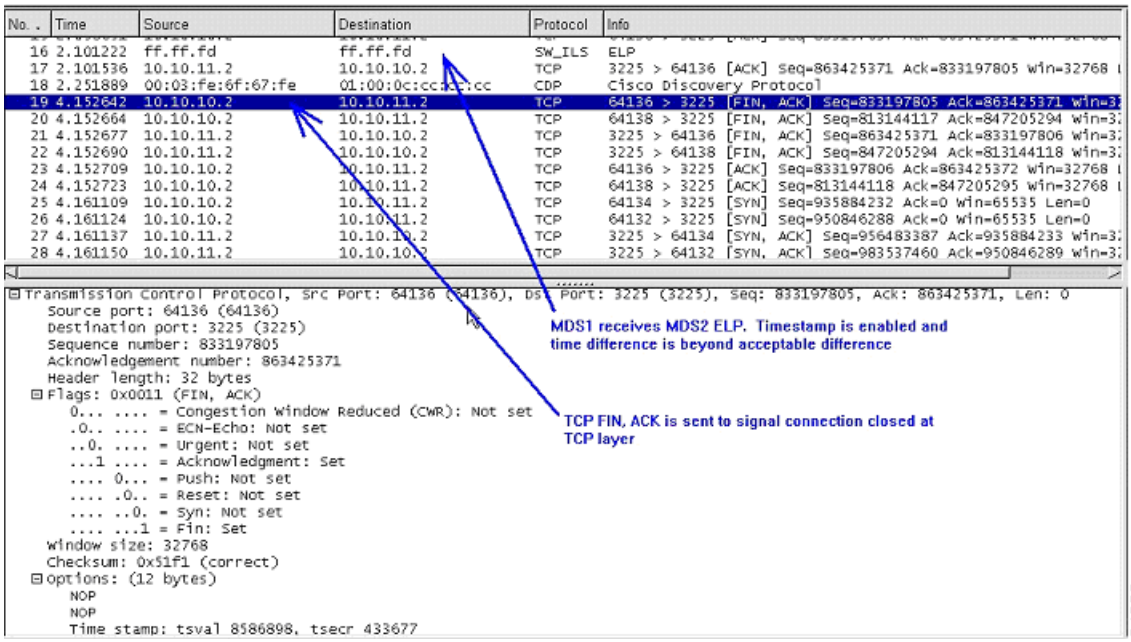

Figure 10-8 shows a trace of timestamp difference accepted.

Figure 10-8 Trace of Timestamp Difference Accepted

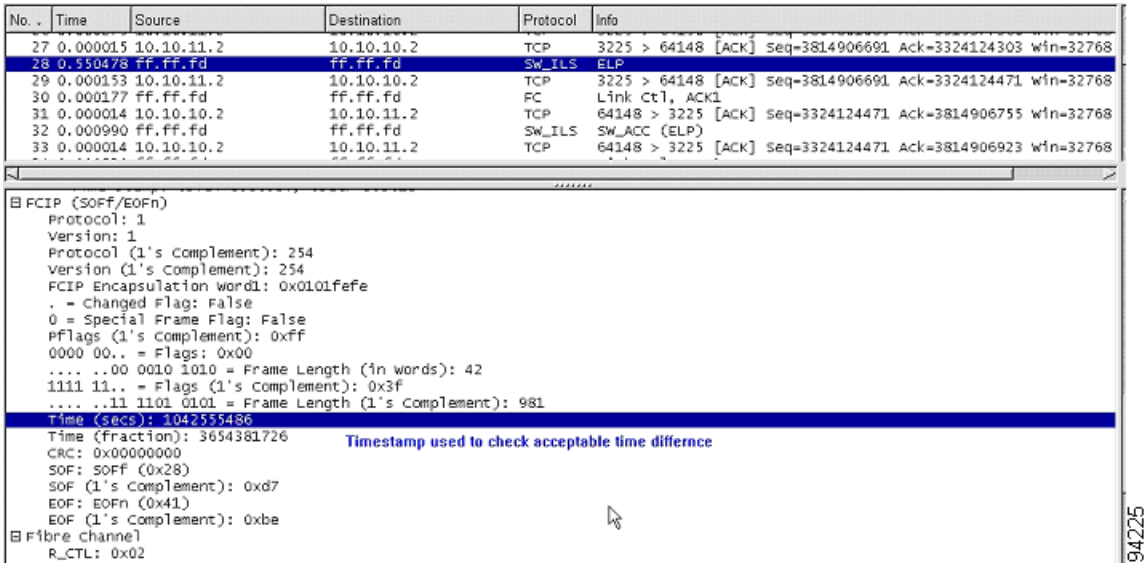

FCIP Special Frame Tunnel Creation and Monitoring

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.

Configuring and Displaying an FCIP Tunnel with a Special Frame with the CLI

<<Add text here.>>

```
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)# interface fcip 21
MDS1(config-if)# special-frame peer-wwn 20:00:00:0b:5f:d5:9f:c0 profile-id 1
```
(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)# interface fcip 21
MDS2(config-if)# special-frame peer-wwn 20:00:00:05:30:00:59:de profile-id 1

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 interface 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 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
```

Figure 10-9 shows a trace of an FCIP tunnel with a Special Frame.

Figure 10-9 Trace of FCIP Tunnel with a Special Frame

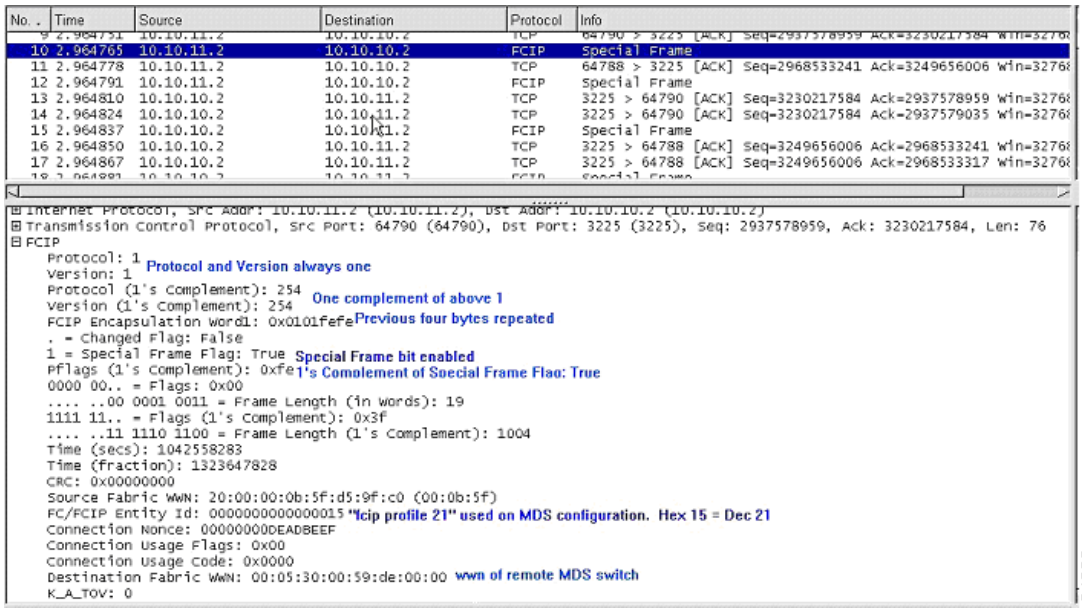

Special Frame Misconfiguration Examples

The following example shows an incorrect peer WWN when using Special Frame.

Displaying Incorrect Peer WWN when Using Special Frame with the CLI

<<Add text here.>>

```
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
```

```
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
```

```
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 10-10 shows a trace of an incorrect remote switch WWN using a Special Frame

Figure 10-10 Trace of Incorrect Remote Switch WWN Using a Special Frame

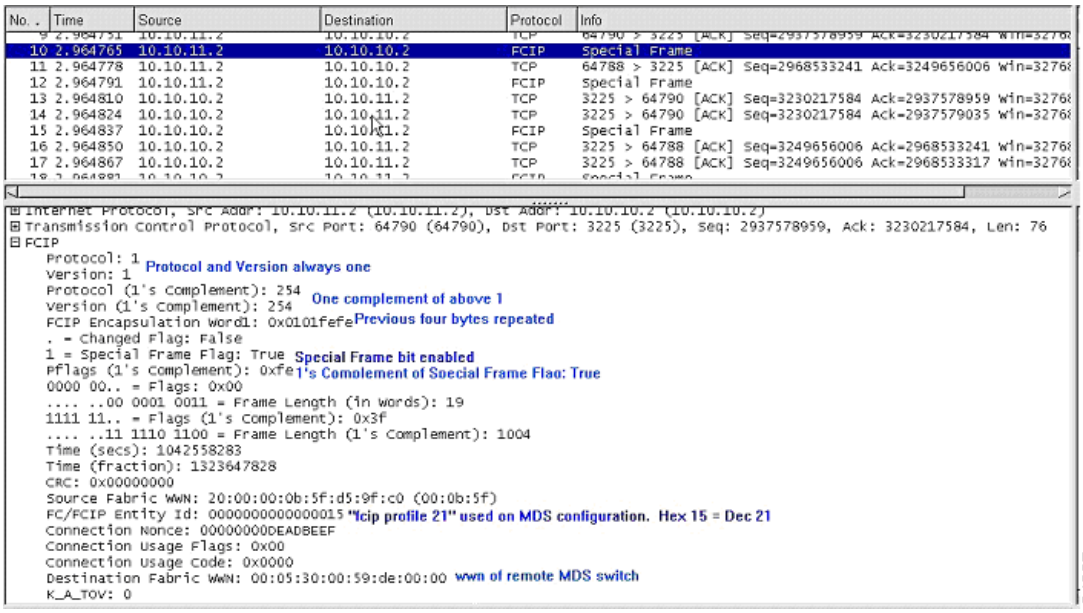

Troubleshooting iSCSI Issues

There are several types of issues you can experience with iSCSI, including the following:

- Troubleshooting iSCSI Authentication, page 10-31
- Displaying iSCSI Authentication with the CLI, page 10-33
- Username/Password Configuration Troubleshooting, page 10-33
- RADIUS Configuration Troubleshooting, page 10-33
- Troubleshooting RADIUS Routing Configuration, page 10-36
- Troubleshooting Dynamic iSCSI Configuration, page 10-36

Troubleshooting iSCSI Authentication

iSCSI user login authentication is required with the Cisco MDS 9000 Family switch. There are two ways of the getting iSCSI users authenticated: either locally configured the in the switch's configuration file, or using the RADIUS server database.

Figure 10-11 shows a successful iSCSI login for the Windows 2000 driver.

Figure 10-11 Successful iSCSI Login Status Window

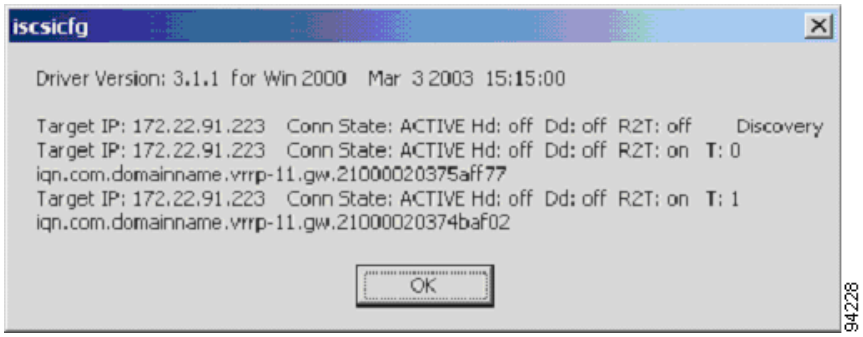

On Solaris systems, a successful login is found in the /var/adm/messages directory, and should look similar to the following example:

```
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
```

Figure 10-12 shows a failed iSCSI login for the Windows 2000 driver.

Figure 10-12 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
```

Displaying iSCSI Authentication with the CLI

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:

```
switch# show authentication
authentication method:none
         console:not enabled
         telnet/ssh:not enabled
authentication method:radius
         console:not enabled
         telnet/ssh:not enabled
         iscsi:not enabled
authentication method:local  <<<<<<<<<<<<<<<<<<<<
         console:enabled
         telnet/ssh:enabled
         iscsi:enabled       <<<<<<<<<<<<<<<<<<<<
switch#
```

If iSCSI is configured for radius authentication, it should looks like this:

```
switch# show authentication
authentication method:none
         console:not enabled
         telnet/ssh:not enabled
authentication method:radius  <<<<<<<<<<<<<<<<<<<<
         console:not enabled
         telnet/ssh:not enabled
         iscsi:enabled        <<<<<<<<<<<<<<<<<<<<
authentication method:local
         console:enabled
         telnet/ssh:enabled
         iscsi:enabled
switch#
```

Username/Password Configuration Troubleshooting

The client side username and password should be check against either the switch's local configuration file or the RADIUS user database.

Verifying iSCSI Users Account Configuration with the CLI

Use the **show user-account** command to verify that the iSCSI users are configured correctly with the username and password, if authentication is against the switch's local user database. Note that the iSCSI password must be at least 16 characters.

```
switch# show user-account iscsi
username:iscsi
secret:1234567812345678

username:iscsiuser
secret:1234567812345678
```

RADIUS Configuration Troubleshooting

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.

Verifying Matching RADIUS Key and Port for Authentication and Accounting with the CLI

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 second and 1 time.

Figure 10-13 shows a Windows-based Radius server configuration.

Figure 10-13 Windows-Based Radius Server Configuration Dialog

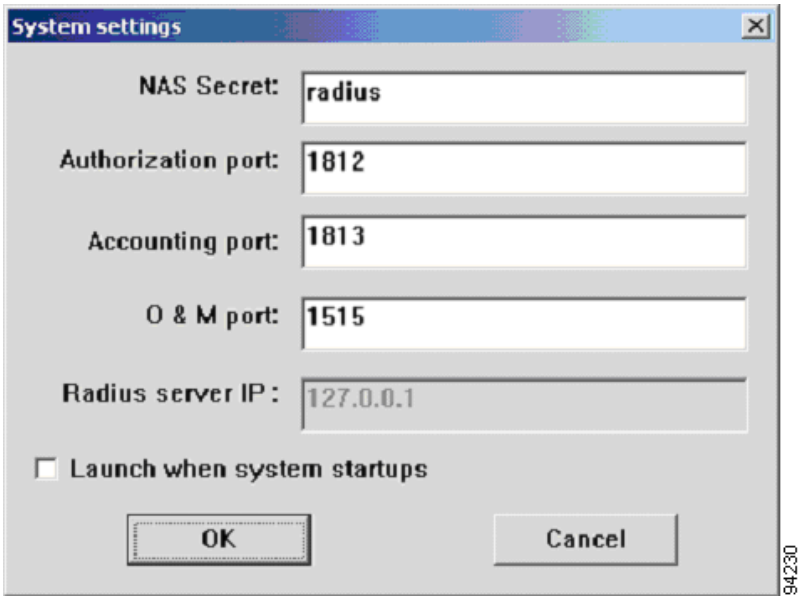

If the items shown above match, verify that the client username and password match those in the Radius database.

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
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÷<¸W¤¤ë-K-ëÙ<]
645 2003y3m14d 15h12m48s CHAP challenge=n8NÝg¤§"__Ó4}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=_¿Éò_à!_AëC0__`õ
656 2003y3m14d 15h12m54s CHAP challenge=_/Ô½Ý×!â£È 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=--5Àùrfàxh
667 2003y3m14d 15h12m54s CHAP challenge=#ùÊÝü{_"__"´_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.
```

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.

Displaying the Debug Output for RADIUS Authentication Request Routing with the CLI

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:
  - With the CLI, invoke the following command:

```
interface Gigabit Ethernet 3/1
no shutdown
.
.
.
interface iscsi 3/1
no shutdown
```

  - With the Fabric Manager, <<Add text here>>

- 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.
  - Do this in Fabric Manager <<Add text here.>>
- 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 Command Output

```
switch# show iscsi session detail
Initiator iqn.1987-05.com.cisco.02.F984BCA7E08C307E2D87A099B2D452F3.FULLMOON (FULLMOON)
  Session #1 (index 2)
    Target iqn.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
      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 Command Output

```
switch# show iscsi remote-node initiator
iSCSI Node name is iqn.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 Command Output

```
switch# show iscsi local-node
target: iqn.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 Command Output

```
switch# show 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 database vsan 1 Command Output

```
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
```

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?
  - How to do it in Fabric Manager <<Add Text here.>>
- 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 with the CLI

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 Command Output

```
switch# show iscsi session detail
Initiator iqn.1987-05.com.cisco.02.8cb3c18879bf356ce18e09679103235f.my-kayak (MY-KAYAK)
  Session #1 (index 84)
    Target iqn.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
      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 #2 (index 85)
    Target iqn.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
    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 iqn.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 iqn.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
```

show iscsi stats Command Output

```
switch# show 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 Command Output

```
switch# show 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:
    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 Command Output

```
switch# show 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 Command Output

```
switch# show 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 Command Output

```
switch# show iscsi remote-node iscsi-session-detail tcp-parameters
iSCSI Node name is iqn.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
      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: iqn.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: iqn.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: iqn.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: iqn.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
```

Fine Tuning/Troubleshooting 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.

Commands Used to Access Performance Data with the CLI

The following CLI commands will give you information related to these criteria.

- **show iscsi remote-node iscsi-session-detail tcp-parameters**
- **show ips stats tcp interface gigabitethernet** *slot*/*port* **detail**
- **show interface iscsi** *slot*/*port*
- **show interface gigabitethernet** *slot*/*port*
- **show interface fc** *slot*/*port*
- **show iscsi remote-node fcp-session-detail**

Understanding TCP Parameters for iSCSI

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 10-14).

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 10-14. The important sections of the displays are highlighted/italicized or bolded.

Figure 10-14 IPS Window Scaling

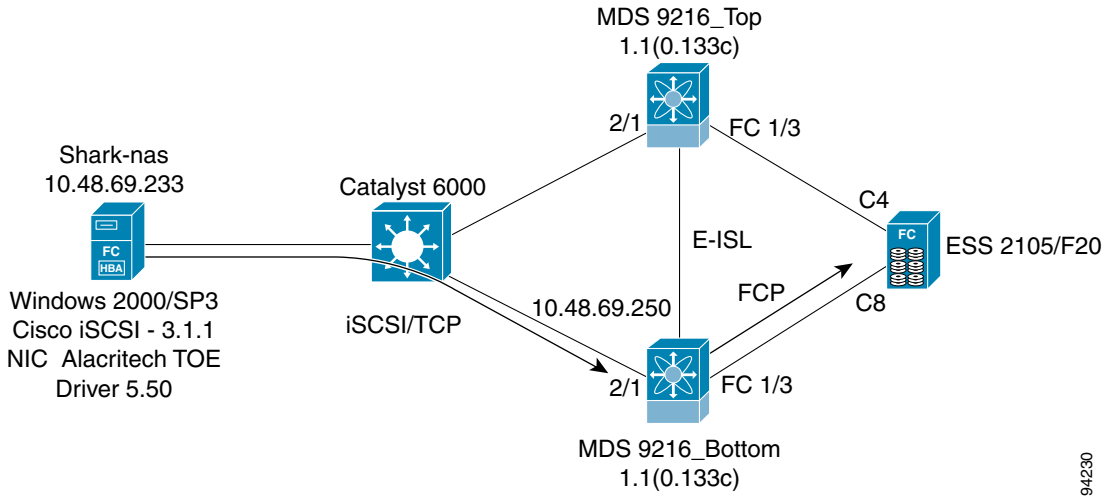

Lab Setup

This is the lab setup that was used in collecting the performance-related information.

The server was an IBM Pentium III 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 WWNs (pWWN) for the IBM Shark storage subsystem. See below for full pWWN:

C4 ➔ 50:05:07:63:00:c4:94:4c (in VSAN 778)

C8 ➔ 50:05:07:63:00:c8:94:4c (in VSAN 777)

Configuring from the Bottom Switch with the CLI

The following example is the configuration for the 9216 switch shown in Figure 10-14.

```
iscsi initiator name iqn.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 iqn.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
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
```

Verifying Connectivity between Client and IPS iSCSI Service

To verify the connectivity between your client and the IPS iSCSI service:

```
MDS_BOTTOM# show ips stats tcp interface gigabitethernet 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)
---  -----------  ------  --------------------------------------------------
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# 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
    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 iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
iscsi tcp
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 20:00:00:0c: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 10-15.)

```
      Peer receive window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4
```

(This is the window size set on the IPS iscsi interface. See Figure 10-15.)

Figure 10-15 Congestion window: Current: 24 KB

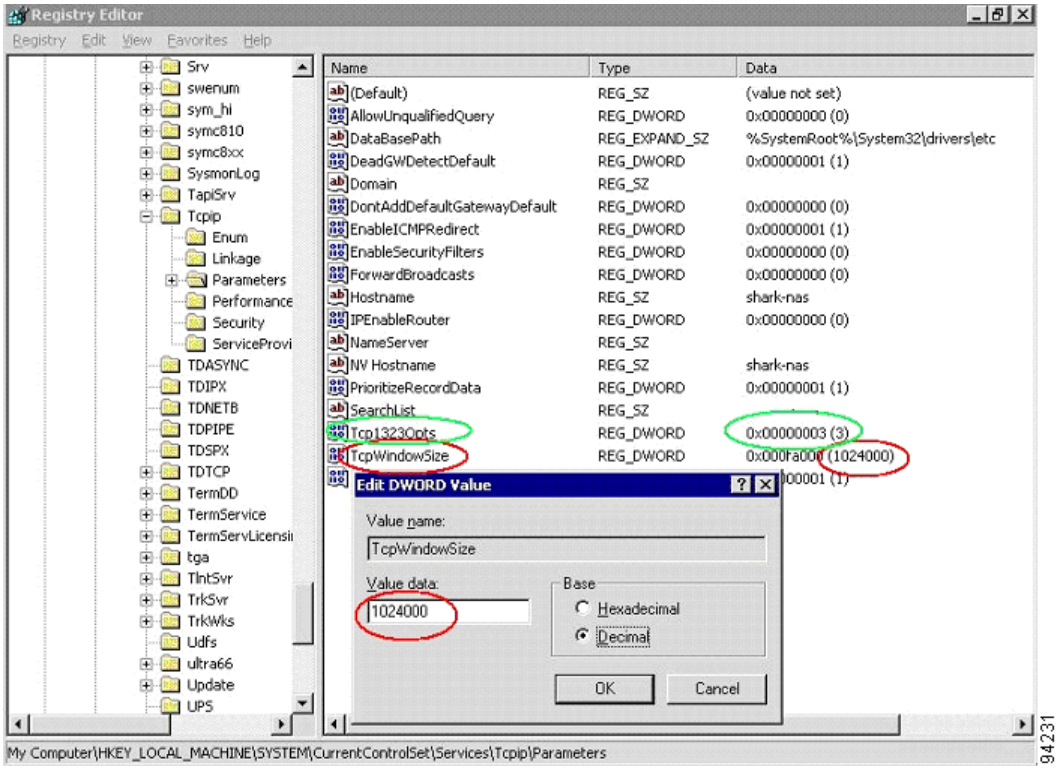

TCP Parameter Changes

To change TCP parameters in the Windows registry, use the Registry Parameters shown in Figure 10-15 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.

Caution Editing the registry is a very high risk operation and can render the system unusable, requiring a reinstallation of the whole operating system. Only advanced users should perform this operation.

Displaying the Gigabit Ethernet Interface with the CLI

<<Add text here.>>

```
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
    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 ACKing 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
      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 unreach, 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_BOTTOM# 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)
                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#
```

Displaying the Effects of Changing the Gigabit MTU on the FC RcvDataFieldSize with the CLI

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
            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
```

Verifying that the Host is Configured for High MTU/MSS with the CLI

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.

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 gigabitethernet 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
```

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 gigabitethernet 2/1 detail
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
...
```

CHAPTER 11

Troubleshooting IPsec

This chapter describes procedures used to troubleshoot IP security (IPsec) and Internet Key Exchange (IKE) encryption in the Cisco MDS 9000 Family Switch products. It includes the following sections:

Overview

The IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is per the latest version of RFC 2401. Cisco MDS SAN-OS IPsec implements RFC 2402 through RFC 2410.

Troubleshooting IPsec Issues

This section provides the procedures required to troubleshoot IKE and IPsec issues in an FCIP configuration. Figure 11-1 shows a simple FCIP configuration where FCIP Tunnel 2 carries encrypted data between switches MDS A and MDS C.

Figure 11-1 Simple FCIP Configuration

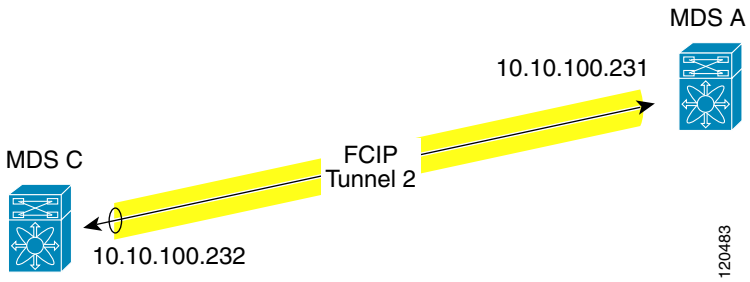

This section includes the following topics:

Verifying IKE Configuration Compatibility

To verify the compatibility of the IKE configurations of MDS A and MDS C shown in Figure 11-1, follow these steps:

Step 1 Ensure that the preshared keys are identical on each switch. Issue the **show crypto ike domain ipsec key** command on both switches. Example command outputs for configuration shown in Figure 11-1 follow:

```
MDSA# show crypto ike domain ipsec key

key ctct address 10.10.100.232

MDSC# show crypto ike domain ipsec key

key ctct address 10.10.100.231
```

Step 2 Ensure that at least one matching policy that has the same encryption algorithm, hash algorithm, and Diffie-Hellman (DH) group, is configured on each switch. Issue the **show crypto ike domain ipsec policy command** on both switches**. Example command outputs** for the configuration shown in Figure 11-1 follow:

```
MDSA# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH group 1

MDSC# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH group 1
```

IPsec Compatibility for iSCSI

Table 11-1 lists the supported settings for encryption and authentication algorithms to be used in iSCSI configurations on the Windows 2000 and Linux platforms.

Table 11-1 Encryption and Authentication Algorithms for iSCSI Configurations

Platform	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Windows 2000 platform	3DES, SHA-1
Cisco iSCSI initiator, FreeSwan IPsec implementation on Linux platform	3DES, MD5

Verifying IPsec Configuration Compatibility

To verify the compatibility of the IPsec configurations of MDS A and MDS C shown in Figure 11-1 follow these steps:

Step 1 Issue the **show crypto map domain ipsec** command and the **show crypto transform-set domain ipsec** command. The following example command outputs display the fields discussed in Step 2 through Step 7.

```
MDSA# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
        Peer = 10.10.100.232
        IP ACL = acl1
            permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
        Transform-sets: tfs-02,
        Security Association Lifetime: 3000 gigabytes/120 seconds
        PFS (Y/N): Y
         PFS Group: group5
Interface using crypto map set cmap-01:
    GigabitEthernet7/1


MDSC# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
        Peer = 10.10.100.231
        IP ACL = acl1
            permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
        Transform-sets: tfs-02,
        Security Association Lifetime: 3000 gigabytes/120 seconds
        PFS (Y/N): Y
         PFS Group: group5
Interface using crypto map set cmap-01:
    GigabitEthernet1/2


MDSA# show crypto transform-set domain ipsec
Transform set:tfs-01 {esp-3des null}
    will negotiate {tunnel}
Transform set:tfs-02 {esp-3des esp-md5-hmac}
    will negotiate {tunnel}
Transform set:ipsec_default_transform_set {esp-aes 128 esp-sha1-hmac}
    will negotiate {tunnel}

MDSC# show crypto transform-set domain ipsec
Transform set:tfs-01 {esp-3des null}
    will negotiate {tunnel}
Transform set:tfs-02 {esp-3des esp-md5-hmac}
    will negotiate {tunnel}
Transform set:ipsec_default_transform_set {esp-aes 128 esp-sha1-hmac}
```

```
    will negotiate {tunnel}
```

Step 2 Ensure the ACLs are compatible on the **show crypto map domain ipsec** command outputs of both switches.

Step 3 Ensure the peer configuration is correct on the **show crypto map domain ipsec** command outputs of both switches.

Step 4 Ensure the transform sets are compatible on the **show crypto transform-set domain ipsec** command outputs of both switches.

Step 5 Ensure that the PFS settings on the **show crypto map domain ipsec** command outputs are configured the same on both switches.

Step 6 Ensure the security association (SA) lifetime settings on the **show crypto map domain ipsec** command outputs are large enough to avoid excessive re-keys (the default settings ensure this).

Step 7 Ensure that the crypto map set is applied to the correct interface on the **show crypto map domain ipsec** command outputs of both switches.

Verifying Security Policy Databases Compatibility

To verify the security policy databases (SPDs) are compatible on both switches, follow these steps:

Step 1 Issue the **show crypto spd domain ipsec** command on both switches to display the SPD. The example command outputs follow:

```
MDSA# show crypto spd domain ipsec
Policy Database for interface:GigabitEthernet7/1, direction:Both
#   0:     deny  udp any port eq 500 any <-----------Clear test policies for IKE
#   1:     deny  udp any any port eq 500 <-----------Clear test policies for IKE
→ #   2:     permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
# 127:     deny  ip any any <------------Clear test policy for all other traffic


MDSC# show crypto spd domain ipsec
Policy Database for interface:GigabitEthernet1/2, direction:Both
#   0:     deny  udp any port eq 500 any
#   1:     deny  udp any any port eq 500
→ #   2:     permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
# 127:     deny  ip any any
```

Step 2 Issue the **show crypto-accelerator interface gigabitethernet** *slot*/*port* **spd inbound** command on both switches to display SPD information from the crypto-accelerator.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

The example command outputs follow:

```
MDSA# show ipsec internal crypto-accelerator interface gigabitethernet 7/1 spd inbound
        Inbound Policy 0 :
        Source IP Address :*
        Destination IP Address :*
        Source port :500, Destination port :* Protocol UDP
        Physical port:0/0, Vlan_id:0/0
        Action cleartext
```

```
        Inbound Policy 1 :
        Source IP Address :*
        Destination IP Address :*
        Source port :*, Destination port :500 Protocol UDP
        Physical port:0/0, Vlan_id:0/0
        Action cleartext

        Inbound Policy 2 :
        Source IP Address :10.10.100.232/255.255.255.255
        Destination IP Address :10.10.100.231/255.255.255.255
        Source port :*, Destination port :* Protocol *
        Physical port:0/1, Vlan_id:0/4095
        Action ipsec

        Inbound Policy 127 :
        Source IP Address :*
        Destination IP Address :*
        Source port :*, Destination port :* Protocol *
        Physical port:0/0, Vlan_id:0/0
        Action cleartext


MDSC# show ipsec internal crypto-accelerator interface gigabitethernet 1/2 spd inbound
        Inbound Policy 0 :
        Source IP Address :*
        Destination IP Address :*
        Source port :500, Destination port :* Protocol UDP
        Physical port:0/0, Vlan_id:0/0
        Action cleartext

        Inbound Policy 1 :
        Source IP Address :*
        Destination IP Address :*
        Source port :*, Destination port :500 Protocol UDP
        Physical port:0/0, Vlan_id:0/0
        Action cleartext

        Inbound Policy 2 :
        Source IP Address :10.10.100.231/255.255.255.255
        Destination IP Address :10.10.100.232/255.255.255.255
        Source port :*, Destination port :* Protocol *
        Physical port:1/1, Vlan_id:0/4095
        Action ipsec

        Inbound Policy 127 :
        Source IP Address :*
        Destination IP Address :*
        Source port :*, Destination port :* Protocol *
        Physical port:0/0, Vlan_id:0/0
        Action cleartext
```

Verifying Interface Status

To verify the status of the interfaces, follow these steps:

Step 1 Issue the **show interface gigabitethernet** command on both switches. Verify that the interfaces are up and their internet addresses are correct. Issue the **no shutdown** command if necessary. The example command outputs follow:

```
MDSA# show interface gigabitethernet 7/1
GigabitEthernet7/1 is up
    Hardware is GigabitEthernet, address is 0005.3001.804e
    Internet address is 10.10.100.231/24
    MTU 1500  bytes
    Port mode is IPS
    Speed is 1 Gbps
    Beacon is turned off
    Auto-Negotiation is turned on
    5 minutes input rate 7728 bits/sec, 966 bytes/sec, 8 frames/sec
    5 minutes output rate 7968 bits/sec, 996 bytes/sec, 8 frames/sec
    7175 packets input, 816924 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    7285 packets output, 840018 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors


MDSC# show interface gigabitethernet 1/2
GigabitEthernet1/2 is up
    Hardware is GigabitEthernet, address is 0005.3001.7f0f
    Internet address is 10.10.100.232/24
    MTU 1500  bytes
    Port mode is IPS
    Speed is 1 Gbps
    Beacon is turned off
    Auto-Negotiation is turned on
    5 minutes input rate 7528 bits/sec, 941 bytes/sec, 8 frames/sec
    5 minutes output rate 7288 bits/sec, 911 bytes/sec, 8 frames/sec
    7209 packets input, 835518 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    7301 packets output, 827630 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Step 2 Issue the **show interface fcip** command on both switches. Verify that each interface is using the correct profile, the peer internet addresses are configured correctly, and the FCIP tunnels are compatible. Issue the **no shutdown** command if necessary. The example command outputs follow:

```
MDSA# show interface fcip 1
fcip1 is trunking
    Hardware is GigabitEthernet
    Port WWN is 21:90:00:0d:ec:02:64:80
    Peer port WWN is 20:14:00:0d:ec:08:5f:c0
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 1 Gbps
    Trunk vsans (admin allowed and active) (1,100,200,302-303,999,3001-3060)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 (100,200,302-303,999,3001-3060)
    Trunk vsans (initializing)             ()
    Using Profile id 1  (interface GigabitEthernet7/1)
    Peer Information
      Peer Internet address is 10.10.100.232 and port is 3225
    FCIP tunnel is protected by IPSec
    Write acceleration mode is off
```

```
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is automatic
    IP Compression is disabled
Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
    TCP Connection Information
      2 Active TCP connections
        Control connection:Local 10.10.100.231:3225, Remote 10.10.100.232:65492
        Data connection:Local 10.10.100.231:3225, Remote 10.10.100.232:65494
      20 Attempts for active connections, 0 close of connections
    TCP Parameters
      Path MTU 1400 bytes
      Current retransmission timeout is 200 ms
      Round trip time:Smoothed 2 ms, Variance:3
      Advertized window:Current:118 KB, Maximum:14 KB, Scale:6
      Peer receive window:Current:128 KB, Maximum:128 KB, Scale:6
      Congestion window:Current:14 KB, Slow start threshold:204 KB
      Current Send Buffer Size:14 KB, Requested Send Buffer Size:0 KB
      CWM Burst Size:50 KB
5 minutes input rate 2960 bits/sec, 370 bytes/sec, 4 frames/sec
    5 minutes output rate 3184 bits/sec, 398 bytes/sec, 4 frames/sec
      3628 frames input, 340644 bytes
         3610 Class F frames input, 338396 bytes
         18 Class 2/3 frames input, 2248 bytes
         0 Reass frames
         0 Error frames timestamp error 0
      3624 frames output, 359140 bytes
         3608 Class F frames output, 357332 bytes
         16 Class 2/3 frames output, 1808 bytes
         0 Error frames


MDSC# show interface fcip 1
fcip1 is trunking
    Hardware is GigabitEthernet
    Port WWN is 20:14:00:0d:ec:08:5f:c0
    Peer port WWN is 21:90:00:0d:ec:02:64:80
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 1 Gbps
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
➔   Using Profile id 1  (interface GigabitEthernet1/2)
    Peer Information
➔     Peer Internet address is 10.10.100.231 and port is 3225
➔   FCIP tunnel is protected by IPSec
    Write acceleration mode is off
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is automatic
    IP Compression is disabled
Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
    TCP Connection Information
```

```
      2 Active TCP connections
        Control connection:Local 10.10.100.232:65492, Remote 10.10.100.231:3225
        Data connection:Local 10.10.100.232:65494, Remote 10.10.100.231:3225
      22 Attempts for active connections, 1 close of connections
    TCP Parameters
      Path MTU 1400 bytes
      Current retransmission timeout is 200 ms
      Round trip time:Smoothed 2 ms, Variance:3
      Advertized window:Current:128 KB, Maximum:14 KB, Scale:6
      Peer receive window:Current:118 KB, Maximum:118 KB, Scale:6
      Congestion window:Current:15 KB, Slow start threshold:204 KB
      Current Send Buffer Size:14 KB, Requested Send Buffer Size:0 KB
      CWM Burst Size:50 KB
5 minutes input rate 3192 bits/sec, 399 bytes/sec, 4 frames/sec
    5 minutes output rate 2960 bits/sec, 370 bytes/sec, 4 frames/sec
      3626 frames input, 359324 bytes
         3610 Class F frames input, 357516 bytes
         16 Class 2/3 frames input, 1808 bytes
         1 Reass frames
         0 Error frames timestamp error 0
      3630 frames output, 340828 bytes
         3612 Class F frames output, 338580 bytes
         18 Class 2/3 frames output, 2248 bytes
         0 Error frames
```

Verifying Security Associations

To verify security associations (SAs), follow these steps:

Step 1 Issue the **show crypto sad domain ipsec** command to verify the current peer, mode, and the inbound and outbound index of each switch. The example command outputs follow:

```
MDSA# show crypto sad domain ipsec
interface:GigabitEthernet7/1
    Crypto map tag:cmap-01, local addr. 10.10.100.231
    protected network:
    local  ident (addr/mask):(10.10.100.231/255.255.255.255)
    remote  ident (addr/mask):(10.10.100.232/255.255.255.255)
➔   current_peer:10.10.100.232
      local crypto endpt.:10.10.100.231, remote crypto endpt.:10.10.100.232
➔     mode:tunnel, crypto algo:esp-3des, auth algo:esp-md5-hmac
      tunnel id is:1
➔    current outbound spi:0x822a202 (136487426), index:1
      lifetimes in seconds::3600
      lifetimes in bytes::483183820800
➔    current inbound spi:0x38147002 (940863490), index:1
      lifetimes in seconds::3600
      lifetimes in bytes::483183820800


MDSC# show crypto sad domain ipsec
interface:GigabitEthernet1/2
    Crypto map tag:cmap-01, local addr. 10.10.100.232
    protected network:
    local  ident (addr/mask):(10.10.100.232/255.255.255.255)
    remote  ident (addr/mask):(10.10.100.231/255.255.255.255)
➔   current_peer:10.10.100.231
      local crypto endpt.:10.10.100.232, remote crypto endpt.:10.10.100.231
```

```
        mode:tunnel, crypto algo:esp-3des, auth algo:esp-md5-hmac
        tunnel id is:1
      current outbound spi:0x38147002 (940863490), index:513
        lifetimes in seconds::3600
        lifetimes in bytes::483183820800
      current inbound spi:0x822a202 (136487426), index:513
        lifetimes in seconds::3600
        lifetimes in bytes::483183820800
```

Step 2 The SA index can be used to look at the SA in the crypto-accelerator. Issue the **show ipsec internal crypto-accelerator interface gigabitethernet** *slot*/*port* **sad** [**inbound** | **outbound**] *sa-index* command to display the inbound or outbound SA information. The hard limit bytes and soft limit bytes fields display the lifetime in bytes. The hard limit expiry secs and the soft limit expiry secs fields display the lifetime in seconds.

Note To issue commands with the **internal** keyword, you must have an account that is a member of the network-admin group.

The example command outputs follow:

```
MDSA# show ipsec internal crypto-accelerator interface gigabitethernet 7/1 sad inbound 1
sw172.22.48.91# show ipsec internal crypto-accelerator interface gigabitethernet 7/1 sad
inbound 1
Inbound SA 1 :
        Mode :Tunnel, flags:0x492300000000000

        IPsec mode is ESP
        Encrypt algorithm is DES/3DES
        Auth algorithm is MD5
        Source ip address 10.10.100.232/255.255.255.255
        Destination ip address 10.10.100.231/255.255.255.255
        Physical port 0, mask:0x1
        Misc select 0 mask:0x0
        Vlan 0 mask:0xfff
        Protocol 0 mask:0x0
        Source port no 0 mask:0x0
        Dest port no 0 mask:0x0
        Hard limit 483183820800 bytes
        Soft limit 401042571264 bytes
        SA byte count 845208 bytes <----Elapsed traffic
        SA user byte count 845208 bytes <----Elapsed traffic
        Error count:auth:0, pad:0, replay:0
        Packet count 7032
        Hard limit expiry 1100652419 secs (since January 1, 1970), remaining 219 7 secs
        Soft limit expiry 1100652386 secs (since January 1, 1970), remaining 216 4 secs
        Sequence number:7033
        Antireplay window:0xffffffff.0xffffffff.0xffffffff.0xffffffff

MDSC# show ipsec internal crypto-accelerator interface gigabitethernet 1/2 sad inbound 513
Inbound SA 513 :
        Mode :Tunnel, flags:0x492300000000000

        IPsec mode is ESP
        Encrypt algorithm is DES/3DES
        Auth algorithm is MD5
        Source ip address 10.10.100.231/255.255.255.255
        Destination ip address 10.10.100.232/255.255.255.255
        Physical port 1, mask:0x1
        Misc select 0 mask:0x0
        Vlan 0 mask:0xfff
        Protocol 0 mask:0x0
```

```
        Source port no 0 mask:0x0
        Dest port no 0 mask:0x0
        Hard limit 483183820800 bytes
        Soft limit 420369924096 bytes
        SA byte count 873056 bytes <----Elapsed traffic
        SA user byte count 873056 bytes <----Elapsed traffic
        Error count:auth:0, pad:0, replay:0
```

Packet count 7137

```
        Hard limit expiry 1100652419 secs (since January 1, 1970), remaining 214 1 secs
        Soft limit expiry 1100652394 secs (since January 1, 1970), remaining 211 6 secs
        Sequence number:7138
        Antireplay window:0xffffffff.0xffffffff.0xffffffff.0xffffffff
```

```
MDSA# show ipsec internal crypto-accelerator interface gigabitethernet 7/1 sad outbound 1
Outbound SA 1 :
        SPI 136487426 (0x822a202), MTU 1400, MTU_delta 4
        Mode :Tunnel, flags:0x92100000000000
        IPsec mode is ESP
        Tunnel options index:0, ttl:0x40, flags:0x1
        Encrypt algorithm is DES/3DES
        Auth algorithm is MD5
        Tunnel source ip address 10.10.100.231
        Tunnel destination ip address 10.10.100.232
        Hard limit 483183820800 bytes
        Soft limit 376883380224 bytes
        SA byte count 874544 bytes <----Elapsed traffic
        SA user byte count 874544 bytes <----Elapsed traffic
        Packet count 7150
        Hard limit expiry 1100652419 secs (since January 1, 1970), remaining 208 9 secs
        Soft limit expiry 1100652384 secs (since January 1, 1970), remaining 205 4 secs
        Outbound MAC table index:1

        Sequence number:7151
```

```
MDSC# show ipsec internal crypto-accelerator interface gigabitethernet 1/2 sad outbound
513
Outbound SA 513 :
        SPI 940863490 (0x38147002), MTU 1400, MTU_delta 4
        Mode :Tunnel, flags:0x92100000000000
        IPsec mode is ESP
        Tunnel options index:0, ttl:0x40, flags:0x1
        Encrypt algorithm is DES/3DES
        Auth algorithm is MD5
        Tunnel source ip address 10.10.100.232
        Tunnel destination ip address 10.10.100.231
        Hard limit 483183820800 bytes
        Soft limit 449360953344 bytes
        SA byte count 855648 bytes <----Elapsed traffic
        SA user byte count 855648 bytes <----Elapsed traffic
        Packet count 7122
        Hard limit expiry 1100652419 secs (since January 1, 1970), remaining 206 4 secs
        Soft limit expiry 1100652397 secs (since January 1, 1970), remaining 204 2 secs
        Outbound MAC table index:125

        Sequence number:7123
```

Security Associations Do Not Re-Key

A lifetime counter (in seconds and bytes) is maintained as soon as an SA is created. When the time limit expires, the SA is no longer operational and is automatically renegotiated (re-keyed) if traffic is present. If there is no traffic, the SA will not be re-keyed and the tunnel will go down.

The re-key operation starts when the soft lifetime expires. That happens approximately 20 to 30 seconds before the time-based lifetime expires, or when approximately 10 to 20 percent of the bytes are remaining in the bytes-based lifetime.

To troubleshoot this problem, follow these steps:

Step 1 Verify that traffic was flowing when the soft SA lifetime expired.

Step 2 Verify that the configurations are still compatible.

Clearing Security Associations

To clear a specific SA, obtain the SA index value and issue the **clear crypto sa domain ipsec interface gigabitethernet** *slot/port* **outbound** *sa-index* command.

To obtain the SA index value, issue the **show crypto sad domain ipsec** command.

Debugging the IPsec Process

To get debug messages printed on the console, the following debugs can be enabled for IPsec:

- debug ipsec error for error messages.
- debug ipsec warning for warning messages.
- debug ipsec config for configuration messages.
- debug ipsec flow for SA related messages.

Debugging the IKE Process

The following commands show the internal state of the IKE process:

- show crypto ike domain ipsec initiator
- show crypto ike domain ipsec sa

Obtaining Statistics from the IPsec Process

To obtain statistics from the IPsec process, issue the **show crypto global domain ipsec** command and the **show crypto global domain ipsec interface gigabitethernet** *slot/port* command. The **show crypto global domain ipsec** command output displays statistics for all SAs. Example command output follows:

```
MDSA# show crypto global domain ipsec
IPSec global statistics:
```

```
        Number of crypto map sets:1
        IKE transaction stats:0 num, 64 max
        Inbound SA stats:1 num
        Outbound SA stats:1 num
```

The **show crypto global domain ipsec interface gigabitethernet** *slot*/*port* command output displays interface level statistics. Example command output follows:

```
MDSA# show crypto global domain ipsec interface gigabitethernet 7/1
IPSec interface statistics:
        IKE transaction stats:0 num
        Inbound SA stats:1 num, 512 max
        Outbound SA stats:1 num, 512 max
```

CHAPTER 12

Troubleshooting Fabric Manager Problems

This chapter contains some common issues you may experience while using Cisco Fabric Manager, and provides solutions.

This chapter contains the following sections:

Tips for Troubleshooting Fabric Manager Problems

This section covers the following topics:

Symptom: The Map Shows Two Switches Where Only One Switch Exists

If two switches show on your map, but you only have one switch, it may be that you have two switches in a non-contiguous VSAN with the same domain ID. The Fabric Manager uses the VSAN ID and domain ID to look up a switch, and this can cause the fabric discovery to assign links incorrectly between these errant switches.

The workaround is to verify that all switches use unique domain IDs within the same VSAN in a physically connected fabric. (The fabric configuration checker will do this task.)

Symptom: Red Line Through the Switch

If a red line shows through your switch, this means the Fabric Manger sees something wrong with the switch. Check the **Switch->Inventory** report. A module, fan, or power supply has failed or is offline and plugged in.

Symptom: Dotted Orange Line Through the Switch

If a dotted orange line shows through your switch, this indicates a minor status warning for that switch. Usually it means an issue with one of the modules. The tooltip should display exactly what is wrong. Hold the mouse over the switch to see the tooltip.

Tips for Using Fabric Manager

This section covers the following topics:

Setting the Map Layout So It Stays After Restarting the Fabric Manager

If you have configured the map layout and would like to “freeze” the map so that the objects stay as they are even after you stop Fabric Manager and restart it again, do the following:

Step 1 Right-click on a blank space in the map. You see a pop-up menu.

Step 2 Select **Layout -> Fix All Nodes** from the menu.

Fabric Manager Upgrade Without Losing Map Settings

When you upgrade from one version of Fabric Manager to another, there is a way to prevent the loss of map settings (enclosure names, placement on the map, etc.)

The $HOME/.cisco_mds9000/db directory contains all the discovered fabrics (*.dat) and maps (*.map). These are upgradable between releases 1.1 and 1.2. If you need to clear the fabric cache, you should first export the enclosures to a file to avoid losing them. Everything else aside from enclosures and map coordinates are stored on the switch. The preferences, last opened, and site_ouis.txt format does not change from release to release.

Restrictions When Using Fabric Manager Across FCIP

Fabric Manager will work without any restrictions across an FCIP tunnel as long as the tunnel is up. However, Fabric Manager cannot automatically discover a Cisco SN5428 mgmt 0 IP address in the fabric. For that switch, it will display a red slash through an FCIP device because of a timeout error. It will still see all targets, initiators, and ISLs attached to a Cisco SN5428 (or any other switch) as long as they appear in the name server or FSPF.

To work around this, you can manually enter the IP address in the Switches table, and click Apply. If the community string is correct, the red slash will go away. Even if the community string is incorrect, double-clicking on the Cisco SN5428 will launch the web tool.

Running Cisco Fabric Manager with Network Multiple Interfaces

If your PC has multiple network interfaces (NICs), the four Cisco Fabric Manager applications detect these interfaces automatically (ignoring loopback interfaces). Fabric Manager client and Device Manager detect all interfaces on your PC each time you launch them, and allow you to select one. Fabric Manager Server and Performance Manager detect on initial install, and allows you to select one. You are not prompted again to choose an interface with these two applications.

There may be circumstances where you will want to change the interface you are using. For example:

- If you add an interface after you have installed Fabric Manager Server and/or Performance Manager
- If you decide to use a different interface than the one you initially selected
- If for any reason one of the Cisco Fabric Manager applications did not detect multiple interfaces

See the following sections, depending on which application you want to recognize the interface.

- Specifying an Interface for Fabric Manager Server, page 12-3
- Specifying an Interface for Fabric Manager Client or Device Manager, page 12-4
- Specifying an Interface for Performance Manager, page 12-3

Specifying an Interface for Fabric Manager Server

To specify an interface for Fabric Manager Server, perform the following steps:

Step 1 Go to the .cisco_mds9000 folder.

Step 2 Edit the server.properties file with a text editor.

Step 3 Scroll until you find the line `snmp.localaddress`.

Step 4 If the line is commented, remove the comment character.

Step 5 Set this value to the IP address or interface name of the NIC you want to use.

Step 6 Save the file.

Step 7 Stop and restart Fabric Manager Server.

Specifying an Interface for Performance Manager

To specify an interface for Performance Manager, perform the following steps:

Step 1 Go to the .cisco_mds9000 folder.

Step 2 Edit the PMCollector.conf file with a text editor.

Step 3 Scroll until you find the line `wrapper.java.additional.2=-Dmds.nmsAddress=`.

Step 4 If the line is commented, remove the comment character.

Step 5 Set this value to the IP address or interface name of the NIC you want to use.

Step 6 Save the file.

Step 7 Stop and restart Performance Server.

Specifying an Interface for Fabric Manager Client or Device Manager

To specify an interface for the Fabric Manager Client or Device Manager, perform the following steps:

Step 1 Go to the .cisco_mds9000/bin folder.

Step 2 Edit the DeviceManager.bat file or the FabricManager.bat file.

Step 3 Scroll to the line that begins with `set JVMARGS=`.

Step 4 Add the parameter -Dmds.nmsaddress=ADDRESS, where ADDRESS is the IP address or interface name of the NIC you want to use.

Step 5 Save the file and relaunch Fabric Manager Client or Device Manager.

Configuring a Proxy Server

If your network uses a proxy server for HTTP requests, make sure the Java Web Start Application Manager is properly configured with the IP address of your proxy server.

To configure a proxy server in the Java Web Start Application Manager, follow these steps:

Step 1 Double-click the Java Web Start application manager icon on your Windows desktop, or choose **Program Files > Java Web Start**.

Step 2 Select **File > Preferences** from the Java WebStart Application Manager.

Step 3 Click the **Manual** radio button and enter the IP address of the proxy server in the HTTP Proxy field.

Step 4 Enter the HTTP port number used by your proxy service in the HTTP Port field.

Step 5 Click **OK**.

Clearing Topology Maps

If you have a switch that you have removed from the fabric, there will be a red X through the switch's icon. You can clear this information from the Fabric Manager client, or from the Fabric Manager server (which will clear the information for all clients) without having to reboot the switch.

To clear information from topology maps, follow these steps:

Step 1 In the Map pane, click on the **Refresh Map** icon.

This clears the information from the client.

Step 2 From the Server menu, click **Purge**.

This clears the information from the server.

Note Any devices not currently accessible (may be offline) will be purged.

Using Fabric Manager in a Mixed Software Environment

You can use Fabric Manager version 2.x to manage a mixed fabric of Cisco MDS 9000 Family switches. Certain 2.x feature tabs will be disabled for any switches running a software version that does not support those features.

APPENDIX A

Before Contacting Technical Support

This appendix describes the steps to perform before calling for technical support for any Cisco MDS 9000 Family multilayer director and fabric switch. This appendix includes the following sections:

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

Steps to Perform Before Calling TAC

At some point, you may need to contact your customer support representative or Cisco TAC for some additional assistance. This section outlines the steps that the you should perform prior to contacting your next level of support, as this will reduce the amount of time spent resolving the issue.

Note Do not reload the module or the switch at least until you have completed Step 1 below. Some logs and counters are kept in volatile storage and will not survive a reload.

To prepare for contacting your customer support representative, follow these steps:

Step 1 Collect switch information and configuration. This should be done before and after the issue has been resolved. The following three methods each provide the same information:

a. Select **Tools > Show Tech Support** in Fabric Manager. Fabric Manager can capture switch configuration information from multiple switches simultaneously. The file can be saved on the local PC.

b. Configure your Telnet or SSH application to log the screen output to a text file. Use the **terminal length 0** CLI command and then use the **show tech-support details** CLI command.

c. Use the **tac-pac <filename>** CLI command to redirect the output of the **show tech-support details** CLI command to a file, and then gzip the file.

```
switch# tac-pac bootflash://showtech.switch1
```

If no filename is specified, the file is created as `volatile:show_tech_out.gz`. The file should then be copied from the switch using the procedure outlined in the "Copying Files to or from the Switch" section on page A-3.

Step 2 If an error occurs in Fabric Manager, take a screen shot of the error. In Windows, press **Alt+PrintScreen** to capture the active window, or press only **PrintScreen** to capture the entire desktop. Then paste this into a new **Microsoft Paint** (or similar program) session and save the file.

Step 3 Capture the exact error codes you see in the message logs from either Fabric Manager or the CLI.

a. Select the **Logs** tab in the Map pane in Fabric Manager or choose **Switches > Events** to see the recent list of messages generated.

b. Copy the error from the message log, which can be displayed using either the **show logging log** CLI command or the **show logging last** *number* to view the last lines of the log.

Step 4 Answer the following questions before calling for technical support:

- On which switch, host bus adapter (HBA), or storage port is the problem occurring?
- Which Cisco SAN-OS software, driver versions, operating systems versions and storage device firmware are in your fabric?
- What is the network topology? (In Fabric Manager, go to **Tools > Show Tech Support** and check the **Save Map** check box.)
- Were any changes being made to the environment (zoning, adding modules, upgrades) prior to or at the time of this event?
- Are there other similarly configured devices that could have this problem, but do not?
- Where was this problematic device connected (which MDS switch and interface)?
- When did this problem first occur?
- When did this problem last occur?
- How often does this problem occur?
- How many devices have this problem?

- Were any traces or debug output captured during the problem time? What troubleshooting steps have you attempted? Which, if any, of the following tools were used?
  - FC Analyzer, PAA-2, Ethereal, local or remote SPAN
  - CLI debug commands
  - FC traceroute, FC ping
  - Fabric Manager or Device Manager tools

Step 5 Is your problem related to a software upgrade attempt?

- What was the original Cisco SAN-OS version?
- What is the new Cisco SAN-OS version?
- Did you use Fabric Manager or the CLI to attempt this upgrade?
- Please collect the output from the following commands and forward them to your customer support representative:
  - **show install all status**
  - **show system internal log install**
  - **show system internal log install details**
  - **show log nvram**

Copying Files to or from the Switch

It may be required to move files to or from the switch. These files may include log, configuration, or firmware files.

Copying Files Using Device Manager

To copy the configuration from the switch using Device Manager, follow these steps:

Step 1 Choose **Admin > Copy Configuration**. You see the Copy Configuration dialog box.

Step 2 Set the To field to the server where you want to copy the configuration file to.

Step 3 Set the From field to running or startup configuration.

Step 4 Select the protocol you want to use to copy the file from the switch.

Step 5 Select **Apply** to copy the file.

To copy files to the switch using Device Manager, follow these steps:

Step 1 Choose **Admin > Flash Files**. You see the list of files in the chosen device and partition.

Step 2 Select **Copy** to copy a file. You see the copy file dialog box.

Step 3 select the protocol you want to use to copy the file to the switch.

Step 4 Set the server address and the file that you want to copy.

Step 5 Select **Apply** to copy the file.

Copying Files Using the CLI

The CLI offers a broad range of protocols to use for copying to or from the switch. Note that the switch always acts as a client, such that an ftp/scp/tftp session will always originate from the switch and either push files to an external system or pull files from an external system.

```
File Server: 172.22.36.10
File to be copied to the switch: /etc/hosts
```

The **copy** CLI command supports four transfer protocols and 12 different sources for files.

```
ca-9506# copy ?
    bootflash: Select source filesystem
    core: Select source filesystem
    debug: Select source filesystem
    ftp: Select source filesystem
    licenses Backup license files
    log: Select source filesystem
    modflash: Select source filesystem
    nvram: Select source filesystem
    running-config Copy running configuration to destination
    scp: Select source filesystem
    sftp: Select source filesystem
    slot0: Select source filesystem
    startup-config Copy startup configuration to destination
    system: Select source filesystem
    tftp: Select source filesystem
    volatile: Select source filesystem
```

Use the following syntax to use secure copy (scp) as the transfer mechanism:

```
"scp:[//[username@]server][/path]"
```

To copy `/etc/hosts` from 172.22.36.10 using the user `user1`, where the destination would be `hosts.txt`, use the following command:

```
switch# copy scp://user1@172.22.36.10/etc/hosts bootflash:hosts.txt
user1@172.22.36.10's password:
hosts 100% |****************************| 2035 00:00
```

To back up the startup-configuration to a sftp server, use the following command:

```
switch# copy startup-config sftp://user1@172.22.36.10/MDS/startup-configuration.bak1
Connecting to 172.22.36.10...
User1@172.22.36.10's password:
switch#
```

Tip Backing up the startup-configuration to a server should be done on a daily basis and prior to any changes. A short script could be written to be run on the MDS to perform a save and then backup of the configuration. The script only needs to contain two commands: **copy running-configuration startup-configuration** and then **copy startup-configuration tftp://*server/name***. To execute the script use: **run-script** filename.

Using Core Dumps

Core dumps are available in situations where unknown problems exist. Dumps are sent to a TFTP server or to a Flash card in slot0: of the local switch. You should set up your switch to generate core dumps under the instruction of your customer support representative. Core dumps are decoded by technical support engineers.

Best practice is to set up cores dumps to go to a TFTP server,. Then these core dumps can be e-mailed directly to your customer support representative.

Setting Up Core Dumps Using the CLI

Use the **system cores** CLI command to set up core dumps on your switch.

```
switch# system cores tftp://10.91.51.200/jsmith_cores
switch# show system cores
Cores are transferred to tftp://10.91.51.200/jsmith_cores
```

Note The file name (indicated by jsmith_cores) must exist in the TFTP server directory.

APPENDIX B

Troubleshooting Tools and Methodology

This appendix describes the troubleshooting tools and methodology available for the Cisco MDS 9000 Family multilayer directors and fabric switches. It includes the following sections:

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 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 Troubleshooting Commands

The command-line interface (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 switches support 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.

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 We recommend that you open a second Telnet or SSH session before entering any debug commands. If the debug session overwhelms the current 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]
```

The following is a summary of some of the common debug commands available Cisco SAN-OS:

Table B-1 Debug Commands

Debug command	Purpose
aaa	Enables AAA debugging.
all	Enables all debugging.
biosd	Enables BIOS daemon debugging.
bootvar	Enables bootvar debugging.
callhome	Enables debugging for Call Home.
cdp	Enables CDP debugging.
cfs	Enables Cisco Fabric Services debugging.
cimserver	Enables CIM server debugging.
core	Enables core daemon debugging.
device-alias	Enables device alias debugging.
dstats	Enables delta statistics debugging.
ethport	Enables port debugging.
exceptionlog	Enables exception log debugging.
fc-tunnel	Enables Fibre Channel tunnel debugging.
fc2	Enables FC2 debugging.
fc2d	Enables FC2D debugging.
fcc	Enables Fibre Channel congestion debugging.
fcdomain	Enables fcdomain debugging.
fcfwd	Enables fcfwd debugging.
fcns	Enables Fibre Channel name server debugging.
fcs	Enables Fabric Configuration Server debugging.
fdmi	Enables FDMI debugging.
flogi	Enables fabric login debugging.
fm	Enables feature manager debugging.
fspf	Enables FSPF debugging.
hardware	Enables hardware, kernel loadable module parameter debugging.
idehsd	Enables idehsd manager debugging.
ilc_helper	Enables ilc-helper debugging.
ipacl	Enables IP ACL debugging.
ipconf	Enables IP configuration debugging.
ipfc	Enables IPFC debugging.
klm	Enables kernel loadable module parameter debugging.
license	Enables license debugging.
logfile	Directs the debug command output to a logfile.
module	Enables module manager debugging.

Table B-1 Debug Commands (continued)

Debug command	Purpose
ntp	Enables NTP debugging.
platform	Enables platform manager debugging.
port	Enables port debugging.
port-channel	Enables PortChannel debug.
qos	Enables QOS Manager debugging.
radius	Enables RADIUS debugging.
rib	Enables RIB debugging.
rlir	Enables RLIR debugging.
rscn	Enables RSCN debugging.
scsi-target	Enables scsi target daemon debugging.
security	Enables security and accounting debugging.
snmp	Enables SNMP debugging.
span	Enables SPAN debugging.
svc	Enables SVC debugging.
system	Enables System debugging.
tlport	Enables TL Port debugging.
vni	Enables virtual network interface debugging.
vrrp	Enables VRRP debugging.
vsan	Enables VSAN manager debugging.
wwn	Enables WWN manager debugging.
zone	Enables zone server debugging.

FC Ping and FC Traceroute

Note Use the Fibre Channel ping and Fibre Channel traceroute features to troubleshoot problems with connectivity and path choices. Do not use them to identify or resolve 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 rerouted and sent back to the source. Using ping, you can verify connectivity and latency to a particular destination across an IP routed network.

The traceroute utility 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 access FC ping and FC traceroute from the CLI or from Fabric Manager.

This section contains the following topics:

Using FC Ping

The FC ping tool:

- Checks end-to-end connectivity.
- Uses an pWWN or FCID.

FC ping allows you to ping a Fibre Channel N port or end device. (See Example B-1.) By specifying the FCID or Fibre Channel address, you can send a series of frames to a target N port. Once these frames reach the target device's N 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 Fibre Channel entity in case of positive or negative answers.

The FC Ping feature verifies reachability of a node by checking its end-to-end connectivity.

- Choose **Tools > Ping** to access FC ping using Fabric Manager.
- Invoke the FC ping feature using the CLI by providing the FC ID or the destination port WWN information in the following ways:

```
switch# fcping pwwn 20:00:00:2e:c4:91:d4:54
switch# fcping fcid 0x123abc
```

Example B-1 FC Ping 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
```

Using FC Traceroute

Use the FC Trace feature to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

FC traceroute identifies the path taken on a hop-by-hop basis and includes a timestamp at each hop in both directions. (See Example B-2.) 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.

Choose **Tools > Traceroute** on Fabric Manager or use the **fctrace** CLI command to access this feature.

Use FC Trace by providing the FC ID, the N port, or the NL port WWN of the destination. The frames are routed normally as long as they are forwarded through TE ports. After the frame reaches the edge of the fabric (the F port or FL port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of failure.

The FC Trace feature works only on TE Ports. Make sure that only TE ports exist in the path to the destination. If there is an E Port in the path:

- The FC Trace frame will be dropped by that switch.
- The FC Trace will time out in the originator.
- Path discovery will not start.

Note FC traceroute will only work across EISL links.

Example B-2 fctraceroute Command

```
switch# fctrace fcid 0xef0000 vsan 1
Route present for :  0xef0000
20:00:00:05:30:00:59:de(0xfffcee)
Latency: 0 msec
20:00:00:05:30:00:58:1e(0xfffc6c)
Timestamp Invalid.
20:00:00:05:30:00:59:1e(0xfffcef)
Latency: 0 msec
20:00:00:05:30:00:59:1e(0xfffcef)
Latency: 174860 msec
20:00:00:05:30:00:58:1e(0xfffc6c)
```

Note The values rendered by the FC traceroute process do not reflect the actual latency across the switches. The actual trace value interpretation is shown in the example below.

```
switch# show fcns database vsan 600

VSAN 600
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPEFEATURE
--------------------------------------------------------------------------
0xeb01e8    NL    210000203767f7a2 (Seagate)       scsi-fcptarget
0xec00e4    NL    210000203767f48a (Seagate)       scsi-fcp
0xec00e8    NL    210000203767f507 (Seagate)       scsi-fcp

Total number of entries = 3
switch# fctrace fcid 0xeb01e8 vsan 600
Route present for   0xeb01e8
2000000530007ade(0xfffcee) ---> MDS originating the trace
Latency 0 msec
2000000c30575ec0(0xfffced) --->first hop MDS towards destination FCID
Latency 30820 msec
2000000c306c2440(0xfffceb) --> MDS which connects directly to the traced FCID (0xeb01e8)
Latency 0 msec
2000000c306c2440(0xfffceb) -->idem, but looped around
Latency 0 msec
```

```
2000000c30575ec0(0xfffced) --> first hop MDS on the return path from traced FCID to
originor
switch#
```

Monitoring Processes and CPUs

There are features in both CLI and Device Manager for monitoring switch processes and CPU status and utilization.

This section contains the following topics:

Viewing Running Processes on Device Manager

Choose **Admin > Running Processes** on Device Manager to view information about the processes currently running on a switch. The Running Processes dialog box (See Figure B-1.)

The dialog display includes:

- Process ID
- The name associated with this process
- The sum of all dynamically allocated memory that this process has received from the system; this includes memory that may have been returned to the system
- The amount of CPU time the process has used, in microseconds

Figure B-1 Running Processes Dialog Box

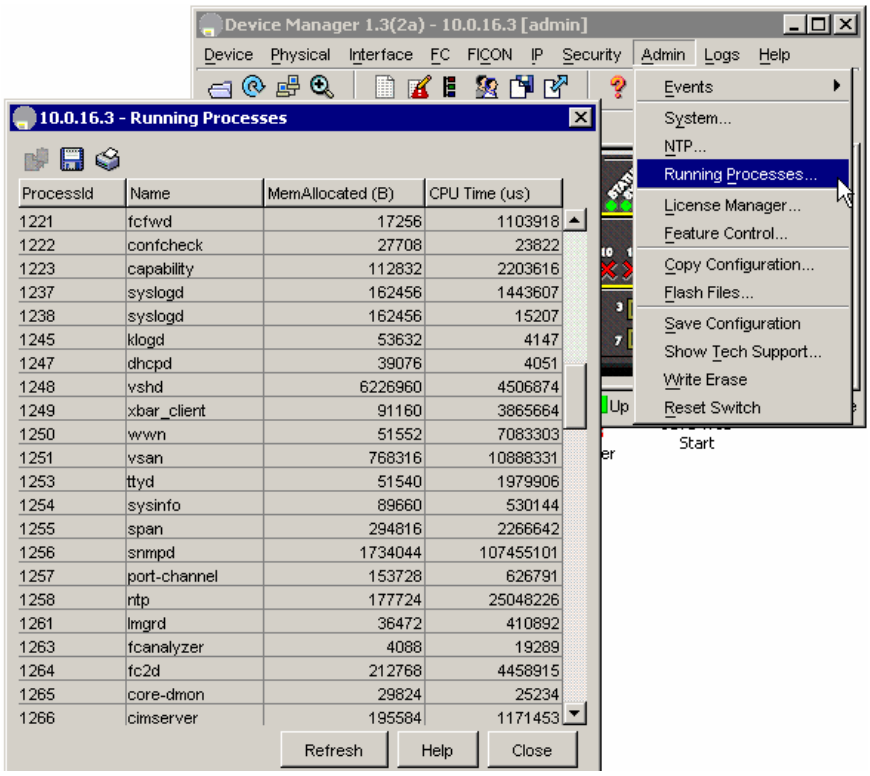

Using the show processes CLI Command

Use the **show processes CLI** command to identify the processes that are running and the status of each process. (See Example B-3.) The command output includes:

- PID = process ID.
- State = process state.
- PC = current program counter in hex format.
- Start_cnt = how many times a process has been started (or restarted).
- TTY = terminal that controls the process. A "-" usually means a daemon not running on any particular TTY.
- Process = name of the process.

Process states are:

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

Note The ER state typically designates a process that has been restarted too many times, causing the system to classify it as faulty and disable it.

Example B-3 show processes Command

```
switch# show processes ?
  cpu     Show processes CPU Info
  log     Show information about process logs
  memory  Show processes Memory Info

switch# show processes
PID    State  PC        Start_cnt    TTY   Process
-----  -----  --------  -----------  ----  -------------
 . . .
  457      S  2abaa76f            1     -  portmap
 1218      S  2acbac24            1     -  licmgr
 1249      S  2ade633e            1     -  xbar_client
 1250      S  2aca833e            1     -  wwn
 1251      S  2aebbc24            1     -  vsan
 1253      S  2ade433e            1     -  ttyd
 1254      S  2ac51ef4            1     -  sysinfo
 1255      S  2af7333e            1     -  span
```

Viewing CPU Time In Device Manager

The Running Processes dialog display can be sorted based on any column header. To sort on CPU utilization, click the CPU column header. An arrow in the column header indicates the order of CPU utilization. Click the column header to toggle between ascending or descending order.

Example B-4 CPU Time Column Header

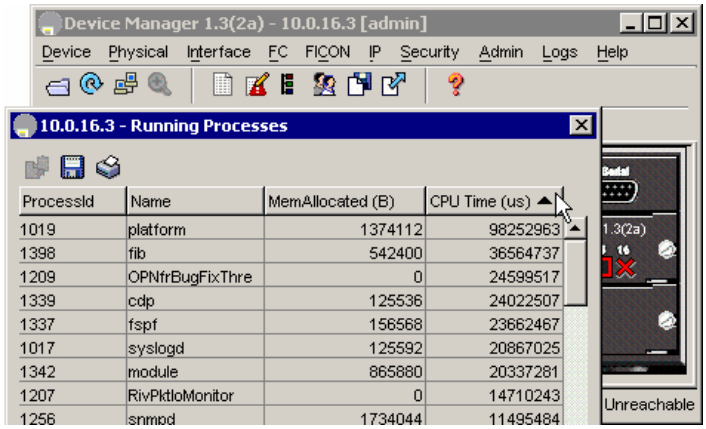

Using the show processes cpu CLI Command

Use the **show processes cpu** CLI command to display CPU utilization. The command output includes:

- Runtime(ms) = CPU time the process has used, expressed in milliseconds.
- Invoked = number of times the process has been invoked.

- uSecs = microseconds of CPU time in average for each process invocation.
- 1Sec = CPU utilization in percentage for the last one second.

Example B-5 show processes cpu Command

```
switch# show processes cpu

PID    Runtime(ms)  Invoked   uSecs  1Sec   Process
-----  -----------  --------  -----  -----  -----------
1016             7         2   3714    0.0  tftpd
1017         20627   2921172      7    0.0  syslogd
1218           299     11710     25    0.0  licmgr
1219            25        38    676    0.0  fs-daemon
1220          1558      6985    223    0.0  feature-mgr
1221           263     11772     22    0.0  fcfwd
1223           512      8996     56    0.0  capability
1237           313     29072     10    0.0  syslogd
1249           912     18815     48    0.0  xbar_client
1250          1481      6214    238    0.0  wwn
1251          1460     68079     21    0.0  vsan
1253           457     29220     15    0.0  ttyd
1254           138      6309     21    0.0  sysinfo
```

Using the show system resource CLI Command

Use the **show system resources** CLI command to display system-related CPU and memory statistics. The output includes the following:

- Load is defined as number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states shows the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage provides the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the used memory statistics.

Example B-6 show system resources Command

```
switch# show system resources

Load average:   1 minute: 0.00   5 minutes: 0.00   15 minutes: 0.00
Processes   :   152 total, 3 running
CPU states  :   0.0% user,   0.0% kernel,   100.0% idle
Memory usage:    960080K total,    412900K used,    547180K free
2340K buffers,  292380K cache
```

Fabric Manager Tools

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:

Note For detailed information about using Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration 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 Tools 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 B-2 shows the Device Manager Summary View window.

Figure B-2 Cisco Device Manager Summary View

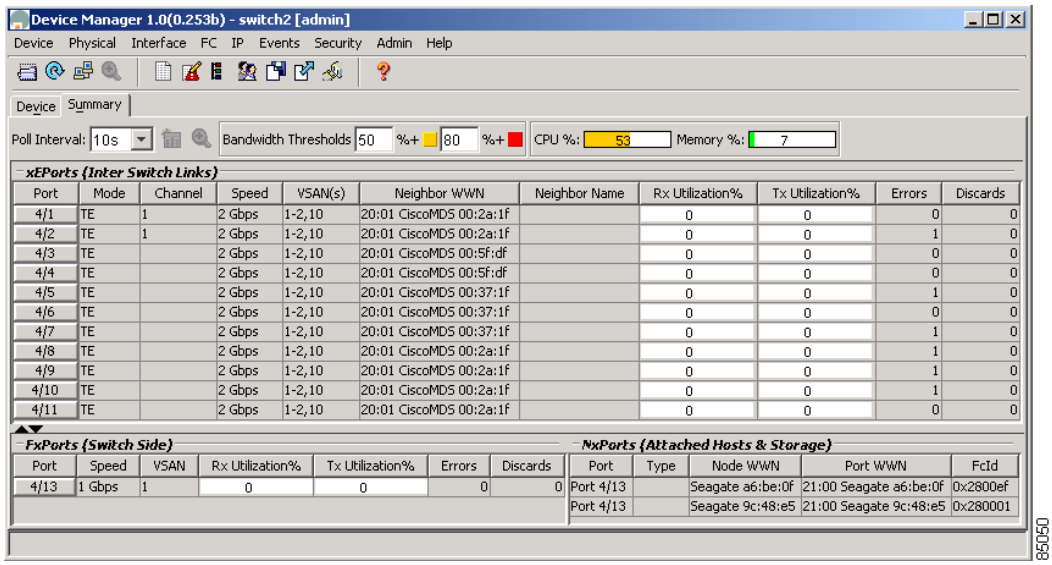

The Summary View window lets you analyze switch performance issues, diagnose problems, and change parameters to resolve problems or inconsistencies. This view shows aggregated statistics for the active Supervisor Module and all switch ports. Information is presented in tabular or graphical formats, with 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 Tools menu to determine the status of the components of a specific switch.

Figure B-3 Switch Health Analysis Window

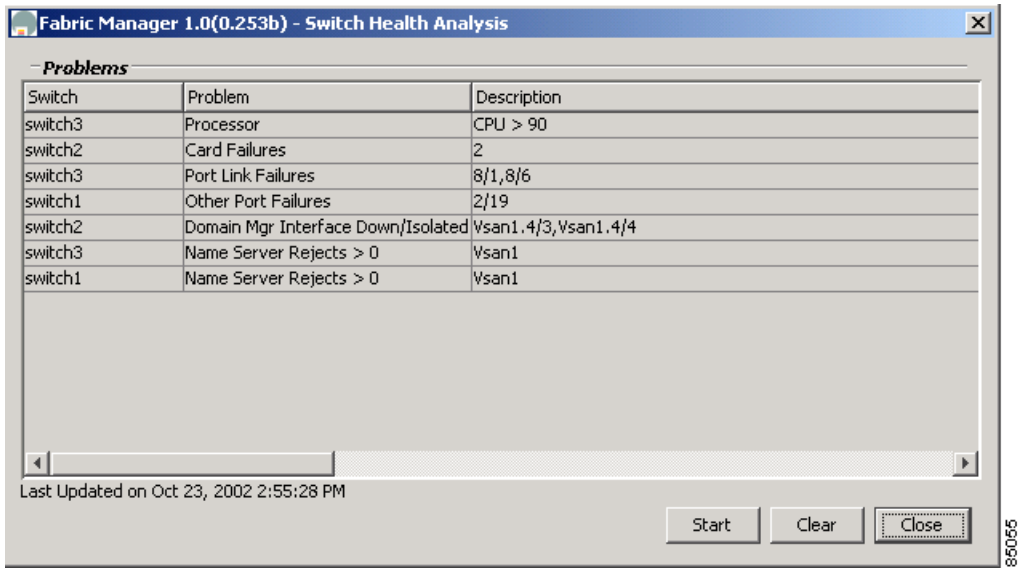

The Switch Health Analysis window displays any problems affecting the selected switches.

Analyzing End-to-End Connectivity

Select **Tools > End to End Connectivity** option from Fabric Manager 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 Tools 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 B-4 Fabric Configuration Analysis Window

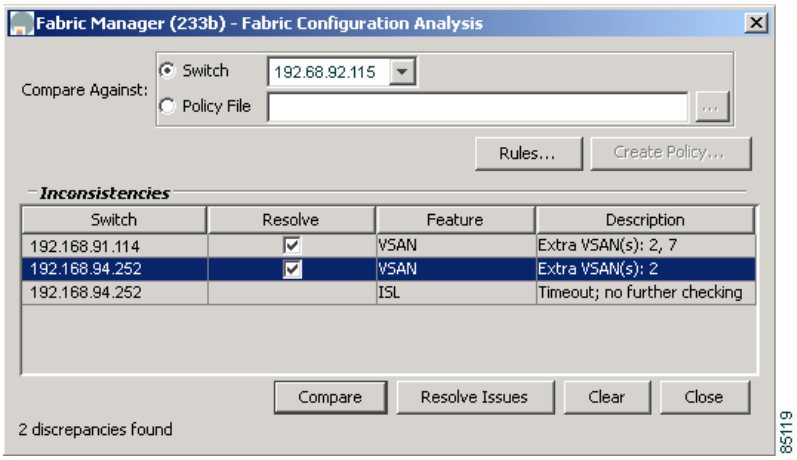

You use a policy file to define the rules to be applied when running the Fabric Checker. When you create a policy file, the system saves the rules selected for the selected switch.

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 Tools menu to determine if two connected switches have compatible zone configurations.

Figure B-5 Zone Merge Analysis Window

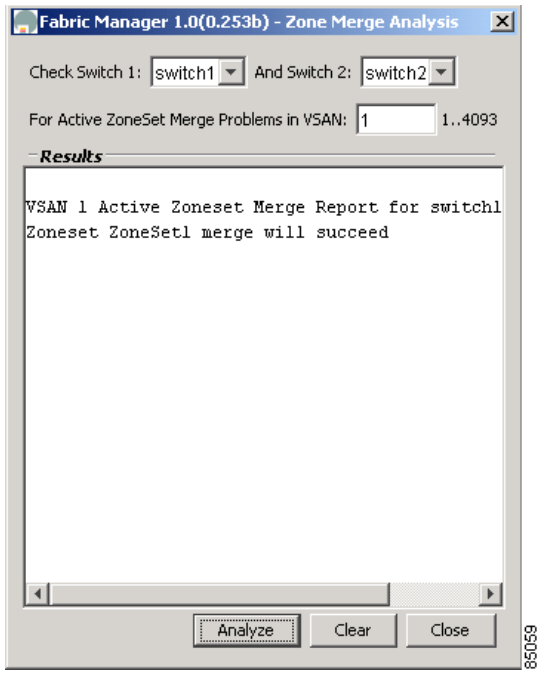

The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches.

You can use the following options on the Fabric Manager Tools 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.

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.

Device Manager: RMON Threshold Manager

Use the options on the Device Manager Events menu to configure and monitor Simple Network Management Protocol (SNMP), Remote Monitor (RMON), Syslog, and Call Home alarms and notifications. SNMP provides a set of preconfigured traps and informs that are automatically generated and sent to the destinations (trap receivers) chosen by the user.

Use the RMON Threshold Manager to configure event thresholds that will trigger log entries or notifications. Use either Fabric Manager or Device Manager to:

- Identify Syslog servers that will record events.
- Configure Call Home, which can issue alerts via e-mail messages or paging when specific events occur.

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, an RMON alarm can be set for a specific level of CPU utilization or crossbar utilization on a switch. The EventGroup allows configuration of events (actions to be taken) based on an alarm condition. Supported event types include logging, SNMP traps, and log-and-trap.

Figure B-6 RMON Threshold Manager

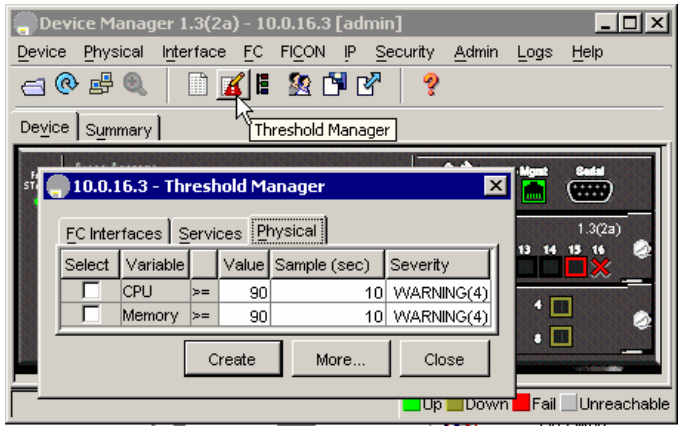

Fibre Channel Name Service

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.

SCSI Target Discovery

The SCSI Target Discovery feature provides 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-targe**t command:

```
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  LR94873000007029 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  LR18591800001004 C:1 A:0 T:3 20:00:00:20:37:18:6f:90
```

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

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 switches provide extensive SNMPv1, v2, and v3support, including Management Information Bases (MIBs) and notifications (traps and informs).

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

SNMPv3 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.

- 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 the *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. See the "Device Manager: RMON Threshold Manager" section on page B-15.

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.

Authorization refers to the scope of access that you have once you have been authenticated. 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
```

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

Using Syslog

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides the following capabilities:

- Logging information for monitoring and troubleshooting.
- Selection of the types of logging information to be captured.
- Selection of the destination of the captured logging information.

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.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM. You can view this log at any time with the **show logging nvram** command.

Logging Levels

The MDS supports the following logging levels:

- 0-emergency
- 1-alert
- 2-critical
- 3-error
- 4-warning
- 5-notification
- 6-informational
- 7-debugging

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. Users can specify which system messages should be saved based on the type of facility and the severity level. Messages are time-stamped to enhance real-time debugging and management.

Enabling Logging for Telnet or SSH

System logging messages are sent to the console based on the default or configured logging facility and severity values.

Users can disable logging to the console or enable logging to a given Telnet or SSH session.

- To disable console logging, use the **no logging console** command in CONFIG mode.
- To enable logging for telnet or SSH, use the **terminal monitor** command in EXEC mode.

Note Note: When logging to a console session is disabled or enabled, that state is applied to all future console sessions. If a user exits and logs in again to a new session, the state is preserved. However, when logging to a Telnet or SSH session is enabled or disabled, that state is applied only to that session. The state is not preserved after the user exits the session.

The **no logging console** command shown in Example B-7:

- Disables console logging
- Enabled by default

Example B-7 no logging console Command

```
switch(config)# no logging console
```

The **terminal monitor** command shown in Example B-8:

- Enables logging for telnet or SSH
- Disabled by default

Example B-8 terminal monitor Command

```
switch# terminal monitor
```

Using Fibre Channel SPAN

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.

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 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 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,
```

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

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 and includes the following topics:

Cisco MDS 9000 Family Port Analyzer Adapter

The Cisco MDS 9000 Family Port Analyzer Adapter is a stand-alone adapter card that converts Fibre Channel 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 MDS 9000 Family Port Analyzer Adapter provides two physical interfaces:

- A Fiber Channel 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 MDS 9000 Family Port Analyzer Adapter does not support half-duplex mode and for this reason, it will not work when connected to a hub.

The Cisco MDS 9000 Family Port Analyzer Adapter provides the following features:

- Encapsulates Fibre Channel frames into Ethernet frames.
- Sustains 32 maximum size Fibre Channel frames burst (in 100 Mbps mode).
- Line rate at 1Gbps (for Fibre Channel frames larger than 91 bytes).
- 64 KB of onboard frame buffer.

- Configurable option for Truncating Fibre Channel frames to 256 bytes (for greater burst).
- Configurable option for Deep Truncating Fibre Channel frames to 64 bytes (best frames burst).
- Configurable option for Ethernet Truncating Fibre Channel frames to 1496 bytes (maximum 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 Fibre Channel frame CRC.

When used in conjunction with the open source protocol analyzer, Ethereal (http//www.ethereal.com), the Cisco MDS 9000 Family 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 MDS 9000 Family Port Analyzer Adapter see the *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Guide*.

Cisco Fabric Analyzer

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 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, 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

fcanalyer 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)
    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.

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

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

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 used to spot broken cables, speed mismatch, link initialization problems and transmission errors. These devices sometimes incorporate higher-level protocol analysis tools and may 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 MDS 9000 Family Port Analyzer Adapter, does not capture and decode at the ordered set level.

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

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

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

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 **perfmeter** 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.

INDEX

Numerics

A

B

C

D

E

F

G

H

T

V

W

Z