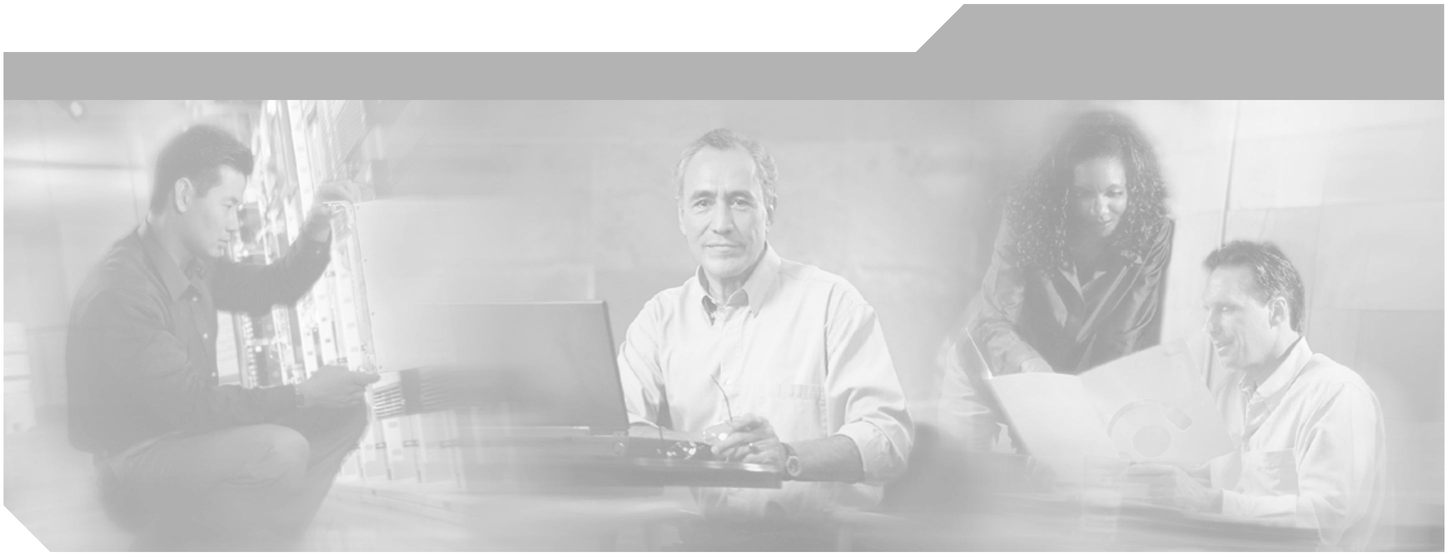

Cisco MDS 9000 Family Cookbook Release 1.x

Cisco MDS SAN-OS Release 1.0(2) through Release 1.3
August 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-8094-01

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

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

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

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

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

Cisco MDS 9000 Family Cookbook

CONTENTS

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

CHAPTER 2

CHAPTER 3

CHAPTER 4

Preface

This preface describes the purpose, audience, organization, and conventions of the *Cisco MDS 9000 Family Cookbook*. It also explains how to obtain related documentation.

Purpose

This document provides step-by-step procedures for configuring and implementing fabrics using the Cisco MDS 9000 Family of multilayer directors and fabric switches. These procedures are intended to complement the procedures in the *Cisco MDS 9000 Family Configuration Guide*, not replace them. Procedures range from simple to complex.The guide also includes several non-technical procedures.

The configurations and components used herein have been tested and validated by Cisco Solution-Interoperability Engineering to support risk-free deployment of fabrics using the Cisco MDS 9000 Family of multilayer directors and fabric switches.

This book is field driven, meaning that the intended audience—storage administrators, technical support engineers, SEs and CEs—is the source for these procedures. Their requirements for a procedure determine the content of this guide. If there are procedures that should be covered in this book, send e-mail to mds-cookbook@cisco.com.

Audience

This document is designed for use by Cisco TAC, sales, support engineers, professional service partners, systems administrators and others who are responsible for the design and deployment of storage area networks in the data center environment.

Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Account Management	Describes how to manage users and their accounts.
Chapter 2	Switch Management	Describes how to perform various configuration and administrative tasks.

Chapter	Title	Description
Chapter 3	Physical Interfaces	Describes how to configure a physical interface.
Chapter 4	Logical Interfaces	Describes how to configure PortChannels.
Chapter 5	VSANs	Describes how to configure and manage VSANs.
Chapter 6	Inter-VSAN Routing	Describes how to configure and manage IVR.
Chapter 7	Zoning	Describes how to configure and manage zones and zone sets.
Chapter 8	iSCSI	Describes how to enable and configure iSCSI.
Chapter 9	FCIP	Describes how to enable and configure FCIP.

Document Conventions

Command descriptions use these conventions:

Convention	Indication
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 }	Required alternative keywords are grouped in braces and separated by vertical bars.
[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:

Convention	Indication
`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.

Tip Means *the following information will help you solve a problem*. These tips are suggested as best practices and are based on in-depth knowledge of the Cisco MDS 9000 family platform and experience implementing SANs.

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

The following warning is specfic to the Cisco 9000 Family of switches.

Warning **Means the following information is critical to ensuring a successful outcome when using the Cisco MDS 9000 Family of switches and Cisco MDS SAN-OS software. Failure to follow the Warning may result in unintended consequences from which it may be difficult to recover.**

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents:

- *Cisco MDS 9000 Family Release Notes for Cisco MDS SAN-OS Releases*
- *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 VERITAS Storage Foundation for Networks Software*
- *Cisco MDS SAN-OS Release Compatibility Matrix for Storage Service Interface Images*
- *Cisco MDS 9000 Family SSM Configuration Note*
- *Cisco MDS 9000 Family ASM Configuration Note*
- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *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*
- *Cisco MDS 9000 Family Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9020 Fabric Switch Configuration Guide and Command Reference*
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Family Fabric and Device Manager Online Help*
- *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide*

- *Cisco MDS 9000 Family Quick Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager Quick Configuration 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*
- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*
- *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Note*

For information on VERITAS Storage Foundation™ for Networks for the Cisco MDS 9000 Family, refer to the VERITAS website: http://support.veritas.com/

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/

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.

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 bug-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

Account Management

This chapter provides recipes for managing users and their accounts. It includes the following sections:

Creating User Accounts for CLI Access

In MDS firmware versions prior to Release 2.0, a separate account was required for both SNMP (Simple Network Management Protocol) and CLI (command-line interface) access. After Release 2.0, a single username grants access to both CLI and SNMP.

Tip

- Use the admin account only during initial setup. After setup, create other user accounts. Each administrator should have their own individual account.
- Always change the admin password from the factory default value.
- Grant users the minimum amount of system rights or privileges required to perform their job.

To access the Cisco MDS 9000 switch via console, SSH, or Telnet, create a username with CLI access. To create a user with CLI access, follow these steps:

Step 1 Enter the configuration submode:

```
switch# config terminal
```

Step 2 Create the CLI user by using the following syntax: username <username> password <password> role <role>

```
switch(config)# username user1 password admin123 role network-admin
```

At this point the user (user1) can access the switch using the password **admin123** via console, SSH or Telnet.

Creating User Accounts for SNMP Access

To access the Cisco MDS 9000 switch using SNMP, create a user with SNMP access. To create a user with SNMP access, follow these steps:

Step 1 Enter the configuration submode:

```
switch# config terminal
```

Step 2 Create the SNMP user by using the following syntax: **snmp-server user <username> <role> auth <encryption method> <password>**

```
switch(config)# snmp-server user user1 network-admin auth md5 admin123
```

At this point, the user (user1) can access the switch using the password **admin123** via an SNMP based product such as Cisco's MDS 9000 Family Fabric Manager or Device Manager.

Tip

- Create a unique user account for each user to aid with troubleshooting and accounting, .
- Use both privacy and authentication passwords for increased security during SNMPv3 based sessions.

Creating an MDS 9000 Switch User Role

The Cisco MDS 9000 switch comes with two predefined roles:

- **Network-admin** is the role assigned to the predefined user called admin. The network-admin canperform any modification to the MDS 9000 platform. There are no restrictions on this user.
- **Network-operator** is a predefined read-only role. The network-operator cannot make modifications to the Cisco MDS 9000 switch. There are no predefined users assigned to this role.

Tip

- Provide each user with a role that provides the minimum number of privileges required to perform their job.
- Leverage the read-only role of the network-operator for those users who do not require the ability to modify the Cisco MDS 9000 switch.
- Use the VSAN based role to allow administrators to have access to and complete control over their VSANs while having read-only or no access to other VSANs.

The following example shows how to create a role that provides the ability to only modify the zoning configuration on the switch.

Step 1 From the **Device Manager**, choose **Security > Common Roles**. (See Figure 1-1.)

Figure 1-1 Common Roles

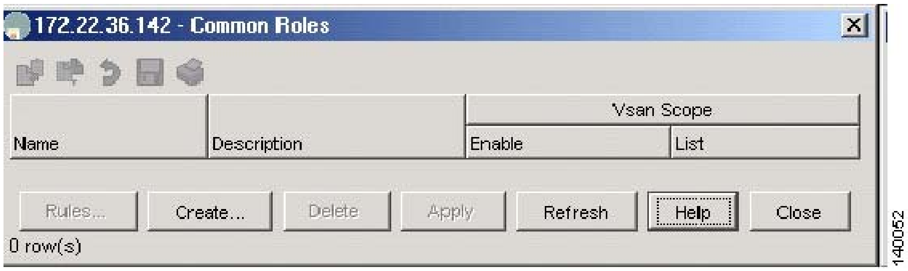

Step 2 Click **Create**.

You see the Create Common Roles dialog box. (See Figure 1-2.)

Figure 1-2 Create Common Roles

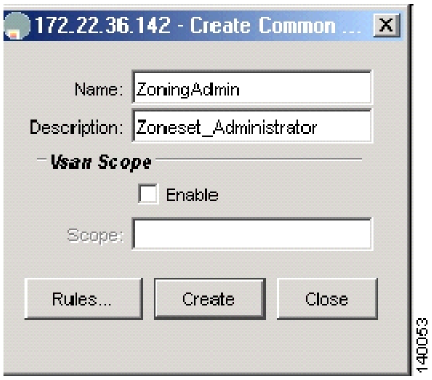

Name: ZoningAdmin

Description: Zoneset_Administrator

□ Enable

Scope:

Step 3 Enter a name and description (without spaces) for the role.

Step 4 Optionally specify a VSAN scope to limit this specific role to a subset of VSANs. A zoning admin role can be created for zone admins who can only modify VSANs 1-10. This example does not specify a VSAN scope.

Step 5 Click **Rules**.

You see the Create Common Role Rules dialog box. (See Figure 1-3.)

Figure 1-3 Create Rules

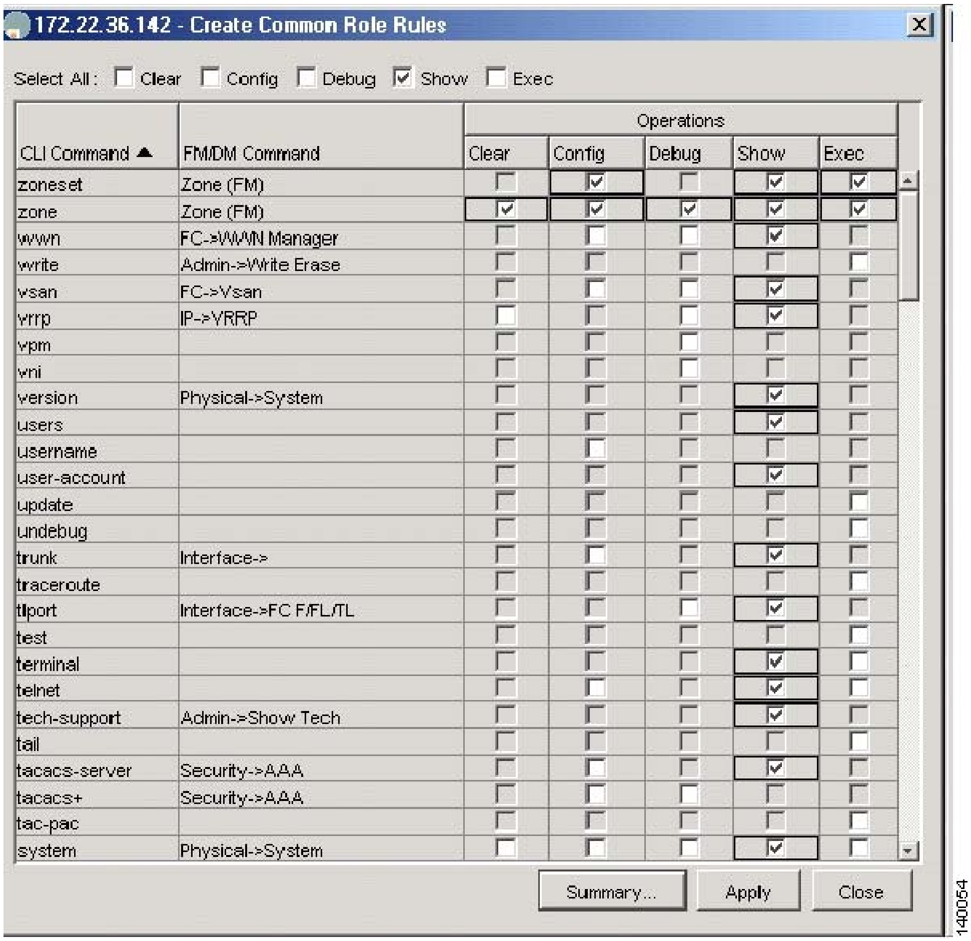

172.22.36.142 - Create Common Role Rules

Select All: ☐ Clear ☐ Config ☐ Debug ☑ Show ☐ Exec

CLI Command ▲	FM/DM Command	Operations				
		Clear	Config	Debug	Show	Exec
zoneset	Zone (FM)	☐	☑	☐	☑	☑
zone	Zone (FM)	☑	☑	☑	☑	☑
wwn	FC->WWN Manager	☐	☐	☐	☑	☐
write	Admin->Write Erase	☐	☐	☐	☐	☐
vsan	FC->Vsan	☐	☐	☐	☑	☐
vrrp	IP->VRRP	☐	☐	☐	☑	☐
vpm		☐	☐	☐	☐	☐
vni		☐	☐	☐	☐	☐
version	Physical->System	☐	☐	☐	☑	☐
users		☐	☐	☐	☑	☐
username		☐	☐	☐	☐	☐
user-account		☐	☐	☐	☑	☐
update		☐	☐	☐	☐	☐
undebug		☐	☐	☐	☐	☐
trunk	Interface->	☐	☐	☐	☑	☐
traceroute		☐	☐	☐	☐	☐
tlport	Interface->FC F/FL/TL	☐	☐	☐	☑	☐
test		☐	☐	☐	☐	☐
terminal		☐	☐	☐	☑	☐
telnet		☐	☐	☐	☑	☐
tech-support	Admin->Show Tech	☐	☐	☐	☑	☐
tail		☐	☐	☐	☐	☐
tacacs-server	Security->AAA	☐	☐	☐	☑	☐
tacacs+	Security->AAA	☐	☐	☐	☐	☐
tac-pac		☐	☐	☐	☐	☐
system	Physical->System	☐	☐	☐	☑	☐

Step 6 Check the **show** check box and all of the operations for the **zone** and **zoneset** commands. Also check the **copy** command so that the zoning admin can save the configuration.

Step 7 Click **Apply**. You see the Create Common Roles dialog box. (See Figure 1-4.)

Figure 1-4 Create Common Roles

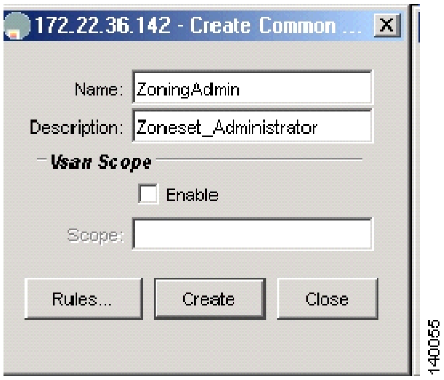

172.22.36.142 - Create Common ...

Name: ZoningAdmin

Description: Zoneset_Administrator

Vsan Scope

☐ Enable

Scope:

Step 8 Click **Create** to display the Display Roles dialog box. (See Figure 1-5.)

Figure 1-5 Display Roles

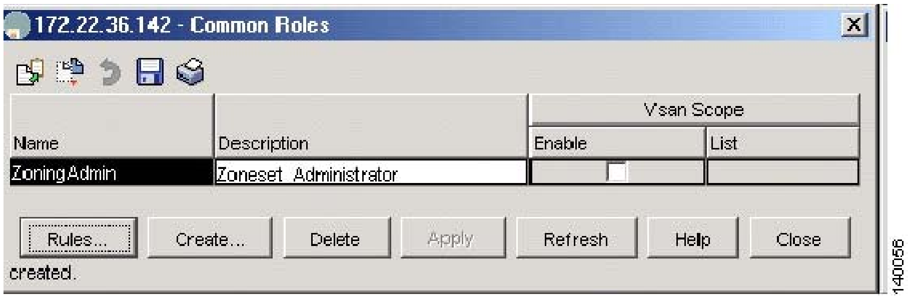

Name	Description	Vsan Scope	
		Enable	List
ZoningAdmin	Zoneset_Administrator	□	

created.

At this point, any user who is assigned to the role of ZoningAdmin can make zoning changes, or use the **copy running-configuration startup-configuration** command.

Step 9 To replicate this role to other systems, examine the startup-configuration for the relevant information, and enter the following CLI commands:

```
switch# config terminal
switch(config)#role name ZoningAdmin
switch(config-role)#  description Zoneset_Administrator
switch(config-role)#  rule 1 permit show
switch(config-role)#  rule 2 permit config feature zoneset
switch(config-role)#  rule 3 permit exec feature zoneset
switch(config-role)#  rule 4 permit clear feature zone
switch(config-role)#  rule 5 permit config feature zone
switch(config-role)#  rule 6 permit debug feature zone
switch(config-role)#  rule 7 permit exec feature zone
switch(config-role)#  rule 8 permit exec feature copy
```

Step 10 To create a user called **zoning_user** with the new role, enter the following commands on the switch:

```
switch# config terminal
switch(config)# username zoning_user password admin123 role ZoningAdmin
```

Configuring TACACS+ with Cisco Secure ACS

Cisco Secure ACS can enhance Cisco MDS 9000 switch management security, and provide centralized authentication, authorization, and accounting for users.

Tip We recommend using a TACACS+ server for both authentication, authorization, and accounting.

Authentication and Authorization with TACACS+

Configuring a Cisco MDS 9000 switch to use TACACS+ allows centralized account management of the switch. Centralized management means that an admin does not have to create and maintain usernames and passwords on individual switches. The Cisco Secure ACS server provides the authentication to a switch login as well as assigns the role to which the user is a member. A shared secret key provides encryption and authentication between the TACACS client (MDS 9500) and the TACACS+ server (Cisco Secure ACS).

In this procedure:

- The switch's IP address is 172.22.36.142
- The TACACS+ server's IP address is 172.22.36.10
- The TACACS+ shared secret key is WarEagle

Configuring Secure ACS Server

Prior to configuring the Cisco MDS 9000 switch, you must configure the Cisco Secure ACS server. To configure Secure ACS to allow the modification of advanced TACACS+ settings, follow these steps:

Step 1 Open Cisco MDS 9000 Family Device Manager. (See Figure 1-6.)

a. In the left pane of the main window, click **Interface Configuration**.

b. In the screen that opens, choose **TACACS+ (Cisco IOS)**.

c. Under **Advanced Configuration Options**, check **Advanced TACACS+ Features** and **Display a window...attributes.**

d. Click **Submit** to save the changes.

Figure 1-6 Secure ACS Configure Display

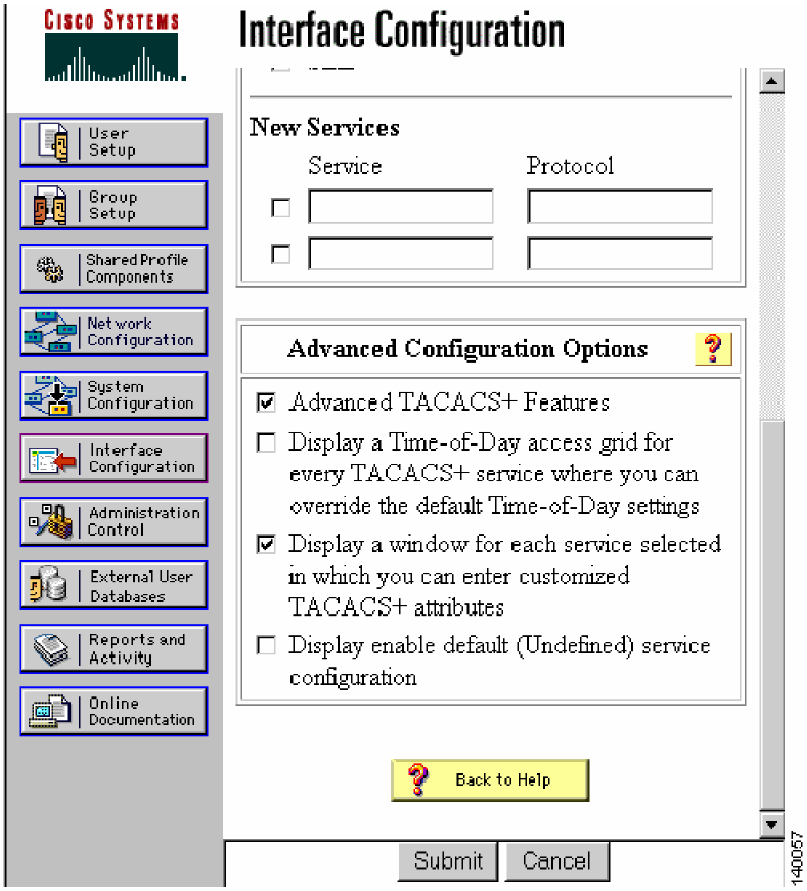

Interface Configuration

New Services

	Service	Protocol
□		
□		

Advanced Configuration Options

☑ Advanced TACACS+ Features

□ Display a Time-of-Day access grid for every TACACS+ service where you can override the default Time-of-Day settings

☑ Display a window for each service selected in which you can enter customized TACACS+ attributes

□ Display enable default (Undefined) service configuration

Step 2 Define the MD 9506 to the TACACS+ server so that the switch can be authenticated by the TACACS+. In the left pane, click **Network Configuration > Add Entry**. (See Figure 1-7.)

a. Enter the MDS 9506 switch's IP address (**172.22.36.142**) and shared secret key (**WarEagle**).

b. Click **Submit** to save the information.

Figure 1-7 Secure ACS Client Setup

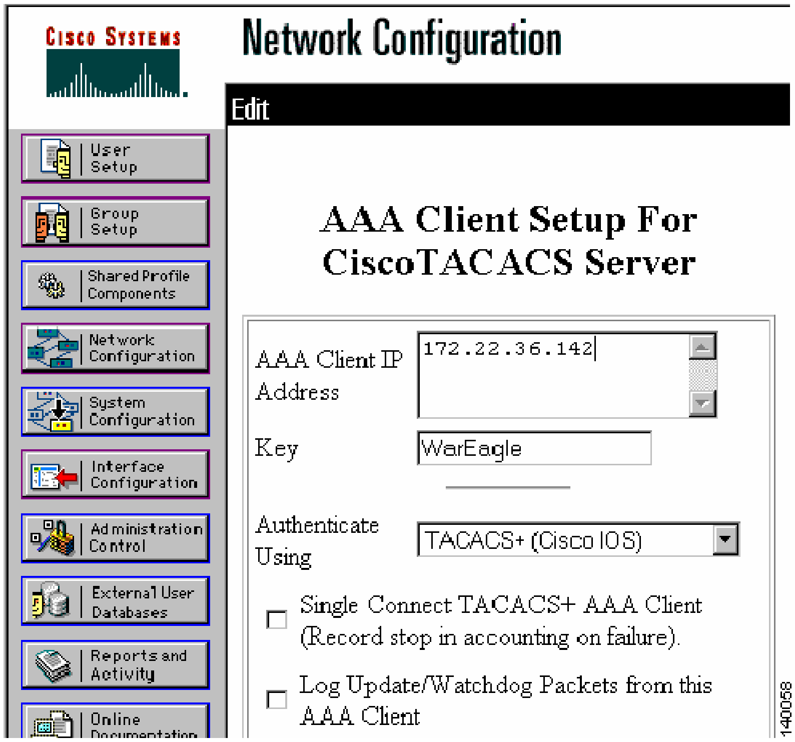

AAA Client Setup For CiscoTACACS Server

AAA Client IP Address: 172.22.36.142

Key: WarEagle

Authenticate Using: TACACS+ (Cisco IOS)

□ Single Connect TACACS+ AAA Client (Record stop in accounting on failure).

□ Log Update/Watchdog Packets from this AAA Client

Step 3 Define a group by clicking **Group Setup**.

Groups provide an easy way to assign the same role to multiple users without having to modify the attributes of each user individually. (See Figure 1-8.)

Figure 1-8 Secure ACS: Group Setup

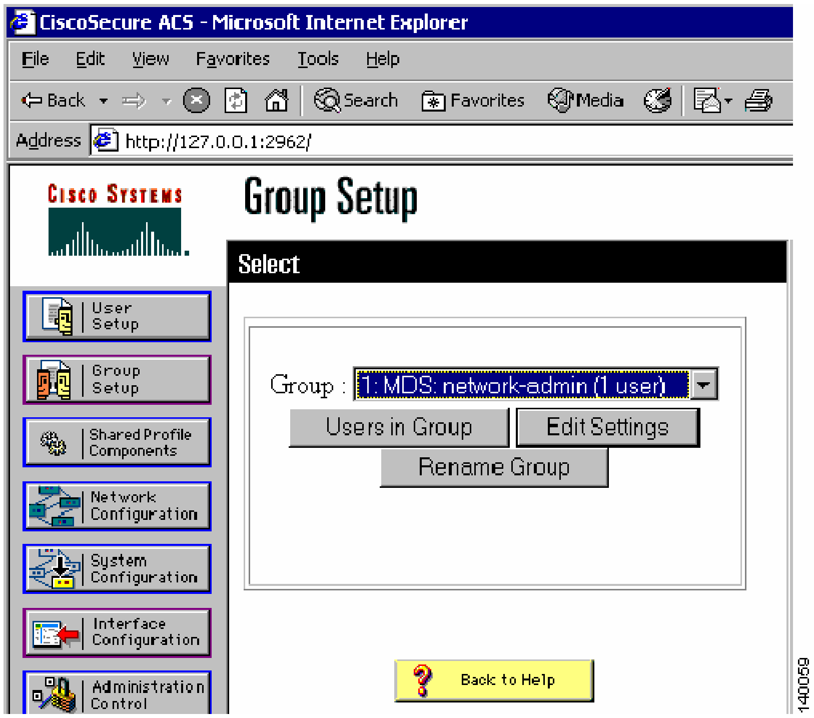

Group : 1: MDS: network-admin (1 user)

Step 4 Chose an available group and click **Rename Group.** In the resulting box, choose a new name for this group.

Tip Use the same Secure ACS group name as the MDS role to ease creation of TACACS+ based users. Click **Submit** to save the name change.

Step 5 In the left pane, click **Group Setup.**

Step 6 Choose the newly renamed group, and click **Edit Settings**.

Step 7 Scroll to the section labeled TACACS+ Settings, check **Shell** and **Custom attributes**.

Step 8 In the Custom attributes field, input the av-pair string corresponding to the role that is defined on the switch for users. The syntax is: `cisco-av-pair=shell:roles="<role>"` Click **Submit + Restart** to save and apply the configuration. (See Figure 1-9.)

Figure 1-9 Secure ACS Adding MDS Role

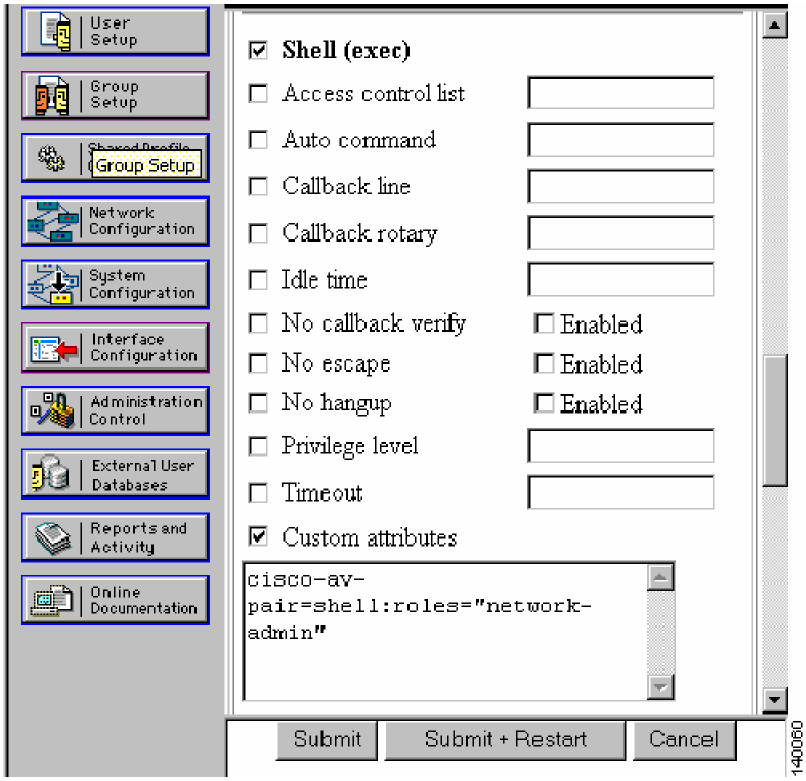

☑ **Shell (exec)**

□ Access control list	
□ Auto command	
□ Callback line	
□ Callback rotary	
□ Idle time	
□ No callback verify	□ Enabled
□ No escape	□ Enabled
□ No hangup	□ Enabled
□ Privilege level	
□ Timeout	

☑ Custom attributes

```
cisco-av-
pair=shell:roles="network-
admin"
```

Step 9 Define a user by clicking **User Setup**.

Step 10 Enter a new or existing username and click **Add/Edit.**

Step 11 In the resulting window, enter a password in the Password and Confirm Password fields.

Step 12 Choose a group from **Group to which the user is assigned** as shown in Figure 1-10.

Figure 1-10 Secure ACS Creating TACACS+ User

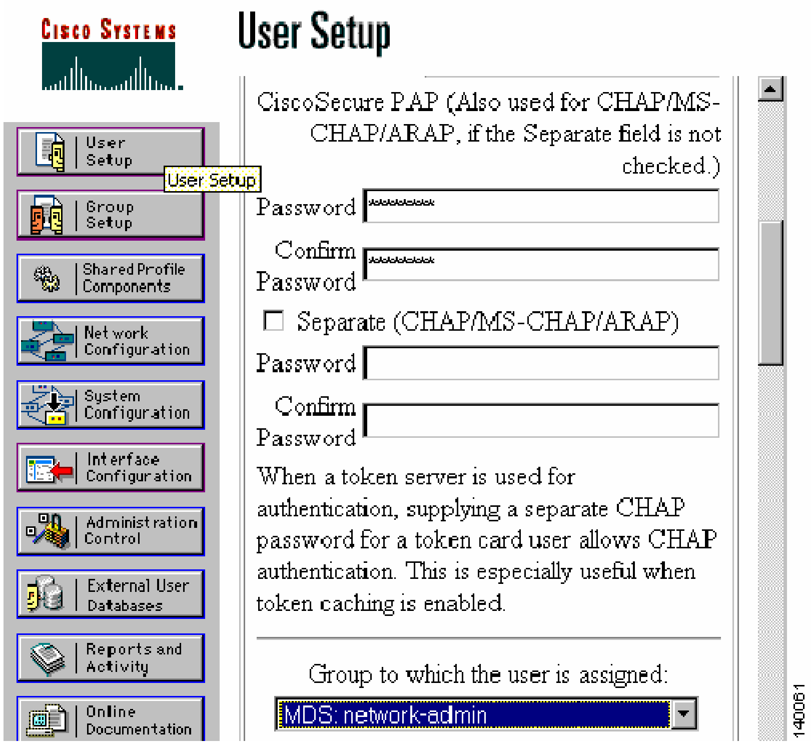

CiscoSecure PAP (Also used for CHAP/MS-CHAP/ARAP, if the Separate field is not checked.)

Password ********

Confirm Password ********

□ Separate (CHAP/MS-CHAP/ARAP)

Password

Confirm Password

When a token server is used for authentication, supplying a separate CHAP password for a token card user allows CHAP authentication. This is especially useful when token caching is enabled.

Group to which the user is assigned: MDS: network-admin

The Secure ACS server configuration is complete. You should now configure the Cisco MDS 9000 switch itself, using either the CLI or the SNMP.

Configuring TACACS+ on the Cisco MDS 9000 Switch

To configure TACACS+ on the switch, follow these steps:

Step 1 Enable TACACS+ by entering the following commands:

```
ca-9506# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# tacacs+ enable
```

Step 2 Define the TACACS+ server and the shared secret key to use with it.

```
ca-9506(config)# tacacs-server host 172.22.36.10 key WarEagle
```

Step 3 Define a group of authentication servers and add the TACACS+ server to the group.

```
ca-9506(config)# aaa group server tacacs+ tacacs-group1
ca-9506(config-tacacs+)# server 172.22.36.10
```

Step 4 Define the methods for the switch to perform authentication for Telnet/SSH/SNMP access.

```
ca-9506(config)# aaa authentication login default group tacacs-group1
```

The following **show** commands display the configuration:

```
ca-9506# show tacacs-server
timeout value:5
total number of servers:1

following TACACS+ servers are configured:
        172.22.36.10:
                available on port:49
                TACACS+ shared secret:********
ca-9506# show aaa authentication
         default: group tacacs-group1
         console: local
         iscsi: local
         dhchap: local

ca-9506# show user-account
user:admin
        this user account has no expiry date
        roles:network-admin

user:seth
        expires on Fri Jun 18 23:59:59 2004
        roles:network-admin
account created through REMOTE authentication
Local login not possible
```

Note The user (seth) is not available locally on the switch, even though seth is a member of the network-admin group or role. This configuration means the user was authenticated by the TACACS+ server and not by the switch.

Accounting with TACACS+

Cisco Secure ACS server can be leveraged to provide a command history that captures which users performed which actions. This information is similar to the CLI **show accounting log** command. However, by placing logging on a remote system, the logs can be independently examined and are available in case the Cisco MDS 9000 switch is inaccessible. This configuration will build upon the configuration defined in Authentication and Authorization with TACACS+, page 1-5.

Configuring the Cisco MDS 9000 Switch

To configure the Cisco MDS 9000 switch to leverage a TACACS+ server for accounting, follow these steps:

Step 1 Configure the Cisco MDS 9000 switch to use the TACACS-group1 server group.

```
switch# conf t
switch(config)# aaa accounting default group  tacacs-group1 local
```

Step 2 Use the **local** keyword to indicate local logging on the switch if all servers listed in the server group are unavailable. If the server group is available, commands or events will **not** be logged locally.

Configuring Cisco Secure ACS

Because this procedure builds on the configuration defined in Authentication and Authorization with TACACS+, page 1-5, only small modifications need to be made.

To configure the Secure ACS server to monitor for Update/Watchdog packets, modify the client configuration by following these steps:

Step 1 In Secure ACS, in the left pane, click **Network Configuration**. (See Figure 1-11.)

Step 2 Choose the client to be modified.

Step 3 Check the **Log Update/Watchdog Packets from this AAA Client** check box.

Figure 1-11 Enabling Accounting on the Secure ACS Server

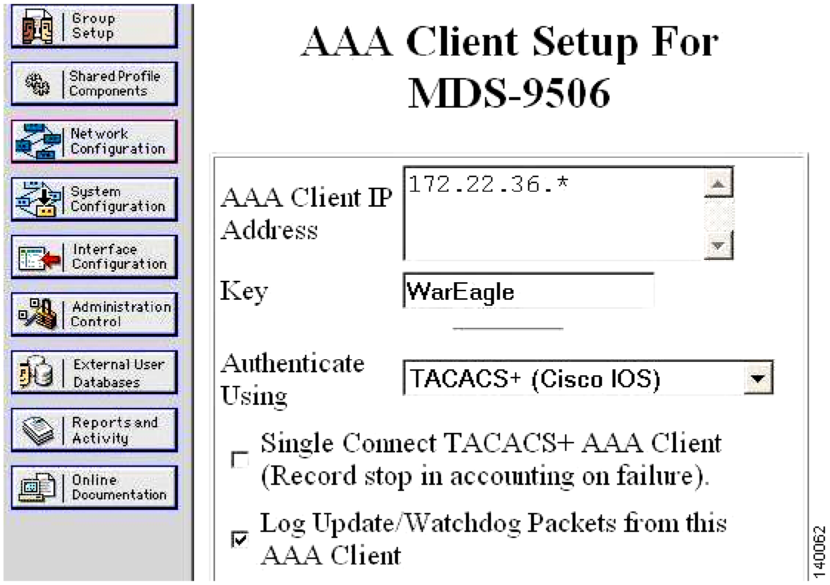

AAA Client Setup For MDS-9506

AAA Client IP Address: 172.22.36.*

Key: WarEagle

Authenticate Using: TACACS+ (Cisco IOS)

☐ Single Connect TACACS+ AAA Client (Record stop in accounting on failure).

☑ Log Update/Watchdog Packets from this AAA Client

To configure the Secure ACS to display commands, follow these steps:

Step 1 Click **System Configuration** in the left pane. (See Figure 1-12.)

Step 2 Click **Logging**.

Step 3 Click **CSV TACACS+ Accounting.**

Step 4 Add the column **err_msg**

Step 5 Check the **Log to CSV TACACS+ Accounting report** box.

Step 6 Click **Submit**.

Figure 1-12 Add MDS Command Logging to Report

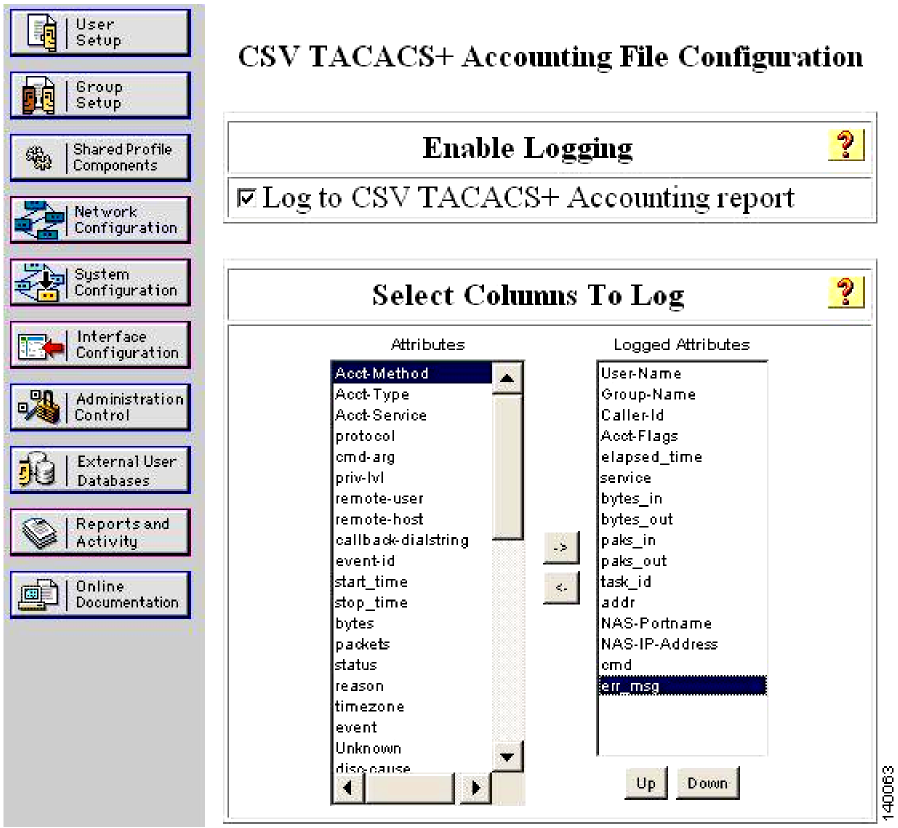

CSV TACACS+ Accounting File Configuration

Enable Logging

☑ Log to CSV TACACS+ Accounting report

Select Columns To Log

Attributes: Acct-Method, Acct-Type, Acct-Service, protocol, cmd-arg, priv-lvl, remote-user, remote-host, callback-dialstring, event-id, start_time, stop_time, bytes, packets, status, reason, timezone, event, Unknown, disc-cause

Logged Attributes: User-Name, Group-Name, Caller-Id, Acct-Flags, elapsed_time, service, bytes_in, bytes_out, paks_in, paks_out, task_id, addr, NAS-Portname, NAS-IP-Address, cmd, err_msg

Step 7 Click **Reports and Activity** in the left pane to view the accounting report. (See Figure 1-13.)

Step 8 Click **TACACS+ Accounting.**

Step 9 In the right pane, choose the day to view. The current day is called **TACACS+ Accounting active.csv**.

Figure 1-13 Secure ACS Accounting Log

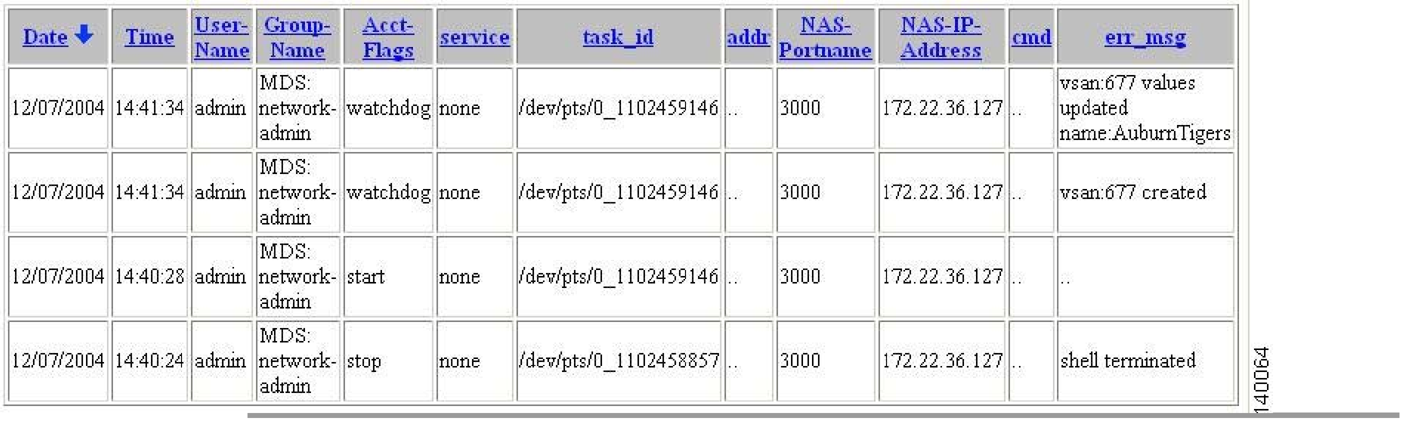

Date ↓	Time	User-Name	Group-Name	Acct-Flags	service	task_id	addr	NAS-Portname	NAS-IP-Address	cmd	err_msg
12/07/2004	14:41:34	admin	MDS: network-admin	watchdog	none	/dev/pts/0_1102459146	..	3000	172.22.36.127	..	vsan:677 values updated name:AuburnTigers
12/07/2004	14:41:34	admin	MDS: network-admin	watchdog	none	/dev/pts/0_1102459146	..	3000	172.22.36.127	..	vsan:677 created
12/07/2004	14:40:28	admin	MDS: network-admin	start	none	/dev/pts/0_1102459146	..	3000	172.22.36.127	..	..
12/07/2004	14:40:24	admin	MDS: network-admin	stop	none	/dev/pts/0_1102458857	..	3000	172.22.36.127	..	shell terminated

Providing Access Without a Password

In some instances, you may need to access the Cisco MDS 9000 switch without using a password, by using automated scripts or agents. Providing a null password or hard coding the password into the script or agent may be considered a weak security practice; however, leveraging the private/public key infrastructure associated with SSH maintains a solid and secure environment.

The procedure includes creating the appropriate key on a host and then adding it to a new user. Because SSH leverages a private/public key exchange, the Cisco MDS 9000 switch knows only the public key, while the host knows both the public and private keys.

Tip Assign logins without passwords to either a read-only role like network-operator or to a role with a minimal set of privileges.

Warning **Having only the public key does not cause the Cisco MDS 9000 switch to grant access to a user; the private key is required to be on the host. The private key should be guarded or treated like a password.**

To set up a read-only (network-operator) based account that only allows access if the user comes from a host that knows both the public and private keys, follow these steps:

Step 1 On the host, create a SSH rsa1 public/private key:

```
$ /usr/bin/ssh-keygen -t rsa1
Generating public/private rsa1 key pair.
Enter file in which to save the key (/users/testuser/.ssh/identity):
/users/setmason/.ssh/identity already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /users/testuser/.ssh/identity.
Your public key has been saved in /users/testuser/.ssh/identity.pub.
The key fingerprint is:
c2:4d:6d:26:21:9d:79:9b:c3:86:dc:a5:07:d2:62:d4 testuser@host
```

On the host, the file /users/testuser/.ssh/identity.pub is the SSH public key that is encrypted using the rsa1 algorithm. The contents of this file are used in the creation of the Cisco MDS 9000 switch user. In this example, the file looks like this:

```
$ cat /users/testuser/.ssh/identity.pub
1024 35
139198677264732164858153476357747926024656548233745027006381178621992083524037906211714241
450436547019604214530354070873624269283640613058470615170649963414635036859628344005142227
886318134122126153182906740418449098047827961768214148936752631482459130056603268404256522
191410368204629699075809390037814979061 testuser@host
```

Step 2 On the Cisco MDS 9000 switch, create all of the SSH keys, even though in this case the client is using rsa1.

```
172.22.36.11# conf t
Enter configuration commands, one per line.  End with CNTL/Z.

172.22.36.11(config)# ssh key rsa1
generating rsa1 key(1024 bits).....
generated rsa1 key
```

```
ca-9506(config)# ssh key dsa
generating dsa key(1024 bits).....
generated dsa key

ca-9506(config)# ssh key rsa
generating rsa key(1024 bits).....
generated rsa key
```

Step 3 Enable SSH on the Cisco MDS 9000 switch.

```
172.22.36.11(config)# ssh server enable
```

Step 4 On the Cisco MDS 9000 switch, create the user by pasting the contents of the identity.pub file after the sshkey parameter:

```
172.22.36.11# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
172.22.36.11(config)# username testuser role network-operator
warning: password for user:setmason not set. S/he cannot login currently
172.22.36.11(config)# username testuser sshkey 1024 35
139198677264732164858153476357747926024656548233745027006381178621992083524037906211714241
450436547019604214530354070873624269283640613058470615170649963414635036859628344005142227
886318134122126153182906740418449098047827961768214148936752631482459130056603268404256522
191410368204629699075809390037814979061 testuser@host
172.22.36.11(config)# end
```

Step 5 Use the following commands to list the user configuration:

```
172.22.36.11# show user-account testuser
user: testuser
        this user account has no expiry date
        roles:network-operator
no password set. Local login not allowed
Remote login through RADIUS/TACACS+ is possible
        ssh public key: 1024 35 139198677264732164858153476357747926024656548233
74502700638117862199208352403790621171424145043654701960421453035407087362426928
36406130584706151706499634146350368596283440051422278863181341221261531829067404
18449098047827961768214148936752631482459130056603268404256522191410368204629699
075809390037814979061 testuser@host
```

Step 6 Test the login process from the host:

```
$ ssh testuser@172.22.36.11
Warning: Remote host denied X11 forwarding.
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2004, 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.
172.22.36.11#
```

If the same user tries logging in from another host that does not have both the private key file (/users/testuser/.ssh/identity) and the public key file (/users/testuser/.ssh/identity), the user will be denied access to the Cisco MDS 9000 switch. The fact that the public key has testuser@host at the end, does not tie it to a specific host, but allows an admin to determine which host it was generated from.

Tip A simple method to leverage this feature is to set up a scheduled backup, for example using cron, to back up the switch configuration on a nightly basis using SSH and TFTP.

Step 7 Set up the backup with the following commands on a host, provided the user has the privileges to issue the **copy** command:

```
#!/bin/sh
#####################################################
#
#/usr/local/bin/backup_mds_config.sh

# This is used for a cron entry. No arguments are
# allowed in cron.Absolute paths to commands must
# be specified to ssh for it to work properly
# ssh key exchange must be separately configured
# for the account "USER"
#
# Adjust the variables for your host and switch
######################################################

DIR=/mds_config
DATE=`date "+%m%d%y_%H%M%S"`
SWITCH_NAME=beat_bama
FILE=$SWITCH_NAME"_run_cfg_"$DATE
USER=cwilliams
COMMAND1="copy running-config startup-config’
COMMAND2="show startup-config"

#Copy running to startup-config
/usr/local/bin/ssh -l $USER $SWITCH_NAME $COMMAND1
#Backup MDS config to local file
/usr/local/bin/ssh -l $USER $SWITCH_NAME $COMMAND2 > $DIR/$FILE
```

The cron job that executes the script must be run by the user specified in the script. Configure the crontab for the user:

```
#Backup MDS config:
00  11  29 4 * /usr/local/bin/backup_mds_config.sh > /tmp/mds_log1
```

CHAPTER 2

Switch Management

This chapter describes various tasks associated with managing a Cisco MDS 9000 switch and includes the following sections:

Saving the Switch Configuration

Saving the configuration after making changes to the Cisco MDS 9000 switch is always a good idea. Whether creating users or configuring ports, the configuration should be saved so that if the switch is rebooted, the current configuration is reapplied to the switch. Optionally, the configuration should also be saved to a file server for purposes of archival, disaster recovery, or version control.

The MDS 9000 switch has two configuration files:

- The running-configuration file describes how the MDS 9000 switch is currently configured.
- The startup-configuration file is the configuration that will be applied to the switch the next time the switch is reloaded.

Both configuration files can be viewed using the **show running-configuration** command or **show startup-configuration** command.

Tip Commands that are listed in the running or startup configuration are valid CLI commands and can be used within the `config terminal` submode on the MDS 9000 switch. Adding `conf t` to the beginning of a file containing CLI commands derived from the running-configuration or the startup-configuration causes the shell to enter the `config` submode.

To save the running-configuration, copy it to the startup-configuration:

```
ca-9506# copy running-config startup-config
[########################################] 100%
```

To copy the startup-configuration to a remote server (in this example the server is SCP), modify the destination filename, by providing a filename to use on the file server (switch1.startupconfig.01182004).

```
ca-9506# copy startup-config scp://user@fileserver/switch1.startup.01182004
setmason@dino's password:
sysmgr_system.cfg    100% |*****************************| 16276       00:00
```

Now the file can be viewed in the switch1.startup.01182004 file.

Copying Files to or from the MDS 9000 Switch

You may need to move files to or from a Cisco MDS 9000 switch. The types of files you may need to move include log files, configuration files, or firmware files. There are two methods for copying files to or from the MDS 9000 switch: using the CLI (command-line interface) and using Fabric Manager.

The first procedure covers the CLI.

The CLI offers a broad range of protocols to use for copying to or from the MDS 9000 switch. Note that the MDS 9000 switch always acts as a client, such that an FTP/SCP/TFTP session always originates from the MDS 9000 switch and either pushes files to an external system or pulls files from an external system.

File Server: 172.22.36.10

File to be copied to the switch: /etc/hosts

The **copy** 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
```

To use SCP (Secure copy) as the transfer mechanism, the syntax is as follows:

scp:[//[username@]server][/path]

To copy /etc/hosts from 172.22.36.10 using user1as the user and the destination filename hosts.txt, enter 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, enter 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 You should back up the startup-configuration to a server on a daily basis and before you make any changes. You can write a short script to run on the MDS 9000 switch to save your configuration and then back it up. The script needs to contain just two commands: **copy running-configuration startup-configuration** and **copy startup-configuration tftp://*server/name***. To execute the script use: **run-script** *filename*.

Managing Files on the Standby Supervisor Module

Occasionally, a file may need to be copied to, copied off, or deleted from the supervisor module, or even deleted from the standby supervisor module. To do this, attach to the standby supervisor module and use the **dir** and **delete** commands.

Note This recipe is most often invoked when a firmware upgrade fails because there is not enough free bootflash: capacity on the standby supervisor for the firmware images.

To perform file copy functions from the supervisor module, follow these steps:

Step 1 Determine which supervisor module is the standby. In this case, it is module 6.

```
switch# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------- ------------------ ------------
1    16     1/2 Gbps FC Module               DS-X9016           ok
2    16     1/2 Gbps FC Module               DS-X9016           ok
3    8      IP Storage Services Module       DS-X9308-SMIP      ok
4    0      Caching Services Module          DS-X9560-SMAP      ok
5    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    active *
6    0      Supervisor/Fabric-1              DS-X9530-SF1-K9    ha-standby
```

Step 2 Connect to the standby supervisor using the **attach** command. Note how the prompt displays the word **standby**.

```
ca-9506# attach module 6
Attaching to module 6 ...
To exit type 'exit', to abort type '$.'
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2004, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or 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.
ca-9506(standby)#
```

Step 3 List the files on the bootflash to be deleted.

```
ca-9506(standby)# dir bootflash:
   12330496     Jun 30 21:11:33 2004  boot-1-3-4a
       2035     Jun 17 16:30:18 2004  hosts.txt
   43705437     Jun 30 21:11:58 2004  isan-1-3-4a
      12288     Dec 31 17:13:48 1979  lost+found/
   12334592     Jun 23 17:02:16 2004  m9500-sf1ek9-kickstart-mz.1.3.4b.bin
   43687917     Jun 23 17:02:42 2004  m9500-sf1ek9-mz.1.3.4b.bin
         99     Apr 07 19:28:54 1980  security_cnv.log

Usage for bootflash://sup-local
  126340096 bytes used
   59745280 bytes free
  186085376 bytes total
```

Step 4 Delete the file with the **delete** command.

```
ca-9506(standby)# delete bootflash:hosts.txt
```

Step 5 To return to the active supervisor, type **exit**. The prompt also returns to the active supervisor prompt.

```
ca-9506(standby)# exit
rlogin: connection closed.
ca-9506#
```

Upgrading MDS 9000 Switch Firmware

To obtain new features and functionality for a Cisco MDS 9000 switch, you may need to upgrade the firmware. You can upgrade using either the CLI or the Fabric Manager.

Firmware images can be downloaded from the Cisco software center located at the following URL: http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml. A CCO login account is required to download all software images.

Tip On single supervisor MDS 9000 switches, like the 9100 and 9200 series, the switch will reboot. Therefore you should enable persistent FC ID and static domain IDs. For information on how to configure these values, see Configuring a Static Domain ID and Persistent FC ID, page 5-4 .

In this procedure the firmware images have been downloaded from the Cisco website and are located on a local file server.

File server: testhost

System image: m9500-sf1ek9-mz.1.3.4b.bin

Kickstart image: m9500-sf1ek9-kickstart-mz.1.3.4b.bin

The location of the firmware images may either be on the switch's bootflash: file system or on another server accessible via FTP/TFTP/SFTP/SCP.

Upgrading Switch Firmware Using the CLI

To upgrade the firmware of an MDS 9000 switch using SCP, enter the following CLI commands:

Step 1 Determine what the upgrade impact will be on the system by using the **show install all impact system** command. This first optional command will also verify the image integrity as well as provide the details of the upgrade. This command does not actually perform the upgrade.

```
ca-9506# show install all impact system
scp://setmason@testhost/tftpboot/rel/qa/1_3_4b/final/m95
00-sf1ek9-mz.1.3.4b.bin kickstart scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-kickstart-mz.1.3.4b.bin
For scp://setmason@testhost, please enter password:
For scp://setmason@testhost, please enter password:

Copying image from scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-kickstart-mz.1.3.4b.bin to
bootflash:///m9500-sf1ek9-kickstart-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Copying image from scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-mz.1.3.4b.bin to
bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-kickstart-mz.1.3.4b.bin
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin
[####################] 100% -- SUCCESS
```

```
Extracting "slc" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "svclc" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:///m9500-sf1ek9-kickstart-mz
.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sf1ek9-kickstart-mz.1.
3.4b.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  non-disruptive       rolling
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset


Other miscellaneous information for installation:
Module  info
------  ----------------------------------


Images will be upgraded according to following table:
Module       Image       Running-Version         New-Version  Upg-Required
------  ----------  --------------------  --------------------  ---------
     1         slc               1.3(4a)               1.3(4b)        yes
     1        bios    v1.1.0(10/24/03)    v1.0.8(08/07/03)         no
     2         slc               1.3(4a)               1.3(4b)        yes
     2        bios    v1.0.8(08/07/03)    v1.0.8(08/07/03)         no
     3         ips               1.3(4a)               1.3(4b)        yes
     3        bios    v1.0.8(08/07/03)    v1.0.8(08/07/03)         no
     4       svclc               1.3(4a)               1.3(4b)        yes
     4       svcsb               1.3(4m)               1.3(4m)         no
     4       svcsb                1.3(4)                1.3(4)         no
     4        bios    v1.1.0(10/24/03)    v1.0.8(08/07/03)         no
     5      system               1.3(4a)               1.3(4b)        yes
     5   kickstart               1.3(4a)               1.3(4b)        yes
     5        bios    v1.1.0(10/24/03)    v1.0.8(08/07/03)         no
     5      loader                1.2(2)                1.2(2)         no
     6      system               1.3(4a)               1.3(4b)        yes
     6   kickstart               1.3(4a)               1.3(4b)        yes
     6        bios    v1.1.0(10/24/03)    v1.0.8(08/07/03)         no
     6      loader                1.2(2)                1.2(2)         no
```

Step 2 Upgrade the firmware using the **install all** command and the appropriate file locations.

```
ca-9506# install all system scp://setmason@testhost/tftpboot/rel/qa/1_3_4b/final/m95
00-sf1ek9-mz.1.3.4b.bin kickstart scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-kickstart-mz.1.3.4b.bin
For scp://setmason@testhost, please enter password:
For scp://setmason@testhost, please enter password:

Copying image from scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-kickstart-mz.1.3.4b.bin to
bootflash:///m9500-sf1ek9-kickstart-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Copying image from scp://setmason@testhost
/tftpboot/rel/qa/1_3_4b/final/m9500-sf1ek9-mz.1.3.4b.bin to
bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-kickstart-mz.1.3.4b.bin
[####################] 100% -- SUCCESS

Verifying image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin
[####################] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "svclc" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:///m9500-sf1ek9-kickstart-mz
.1.3.4b.bin.
[####################] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sf1ek9-kickstart-mz.1.
3.4b.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  non-disruptive       rolling
     4       yes  non-disruptive       rolling
     5       yes  non-disruptive         reset
     6       yes  non-disruptive         reset


Other miscellaneous information for installation:
Module  info
------  ----------------------------------



Images will be upgraded according to following table:
```

```
Module       Image       Running-Version        New-Version  Upg-Required
------  ----------  --------------------  --------------------  ---------
     1         slc               1.3(4a)               1.3(4b)        yes
     1        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)         no
     2         slc               1.3(4a)               1.3(4b)        yes
     2        bios      v1.0.8(08/07/03)      v1.0.8(08/07/03)         no
     3         ips               1.3(4a)               1.3(4b)        yes
     3        bios      v1.0.8(08/07/03)      v1.0.8(08/07/03)         no
     4       svclc               1.3(4a)               1.3(4b)        yes
     4       svcsb               1.3(4m)               1.3(4m)         no
     4       svcsb                1.3(4)                1.3(4)         no
     4        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)         no
     5      system               1.3(4a)               1.3(4b)        yes
     5   kickstart               1.3(4a)               1.3(4b)        yes
     5        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)         no
     5      loader                1.2(2)                1.2(2)         no
     6      system               1.3(4a)               1.3(4b)        yes
     6   kickstart               1.3(4a)               1.3(4b)        yes
     6        bios      v1.1.0(10/24/03)      v1.0.8(08/07/03)         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.1.3.4b.bin to standby.
[####################] 100% -- SUCCESS

Syncing image bootflash:///m9500-sf1ek9-mz.1.3.4b.bin to standby.
[####################] 100% -- SUCCESS

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

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

Module 5: Waiting for module online.
 -- SUCCESS
```

At this point, the switch performs a hitless supervisor switchover. A new Telnet/CLI session must be established to the new supervisor.

Note If the images fail to copy to the standby supervisor, there may be insufficient room for the new images and some old images or files may need to be removed. See Managing Files on the Standby Supervisor Module, page 2-3 for a recipe on removing files from the standby supervisor.

Step 3 To view the status of the current upgrade from the new supervisor, enter the **show install all status** command.

```
switch# show install all status
This is the log of last installation.

Continue on installation process, please wait.
The login will be disabled until the installation is completed.

Module 5: Waiting for module online.
 -- SUCCESS

Module 1: Non-disruptive upgrading.
 -- SUCCESS
```

```
Module 2: Non-disruptive upgrading.
 -- SUCCESS

Module 3: Non-disruptive upgrading.
 -- SUCCESS

Module 4: Non-disruptive upgrading.
 -- SUCCESS

Install has been successful.
```

Upgrading Switch Firmware Using Fabric Manager

To upgrade the firmware of one or more MDS 9000 switches, leverage the interface of the Fabric Manager and follow these steps:

Step 1 Select the **Software Install Wizard** from the toolbar in Fabric Manager. (See Figure 2-1.)

Figure 2-1 Image Installation with Fabric Manager

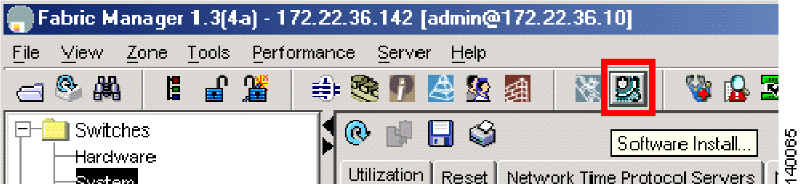
Fabric Manager 1.3(4a) - 172.22.36.142 [admin@172.22.36.10]
File View Zone Tools Performance Server Help
Switches
Hardware
Software Install...
Utilization
Reset
Network Time Protocol Servers
140065

Step 2 Choose the switches to upgrade and click **Next**. (See Figure 2-2.)

Figure 2-2 Choose Switches to Upgrade

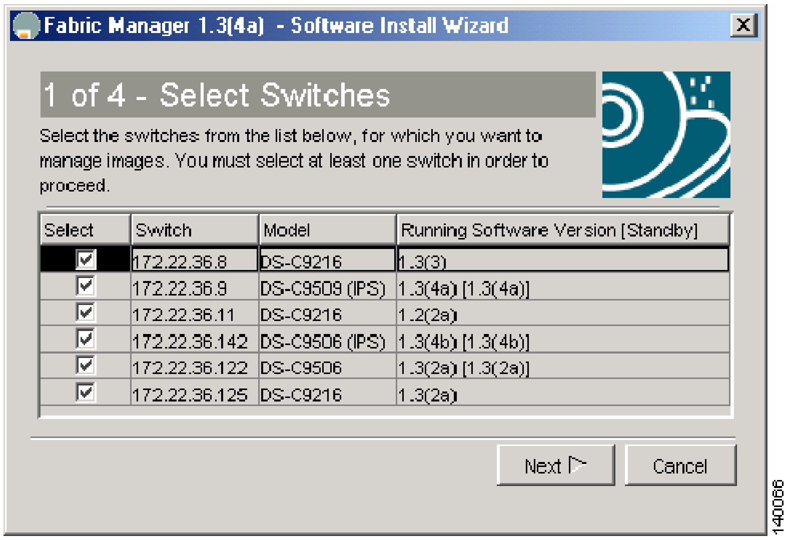

1 of 4 - Select Switches

Select the switches from the list below, for which you want to manage images. You must select at least one switch in order to proceed.

Select	Switch	Model	Running Software Version [Standby]
☑	172.22.36.8	DS-C9216	1.3(3)
☑	172.22.36.9	DS-C9509 (IPS)	1.3(4a) [1.3(4a)]
☑	172.22.36.11	DS-C9216	1.2(2a)
☑	172.22.36.142	DS-C9506 (IPS)	1.3(4b) [1.3(4b)]
☑	172.22.36.122	DS-C9506	1.3(2a) [1.3(2a)]
☑	172.22.36.125	DS-C9216	1.3(2a)

Step 3 Specify the location of the firmware images. (See Figure 2-3.)

a. Enter the file information to transfer the file from the server to the switch.

b. If the files are to be downloaded during the install, also enter the path and filename of the images.

c. By checking the **Skip Image Download** check box, an upgrade can be performed using images that are already located on the supervisor's bootflash.

Figure 2-3 Specify Firmware Images

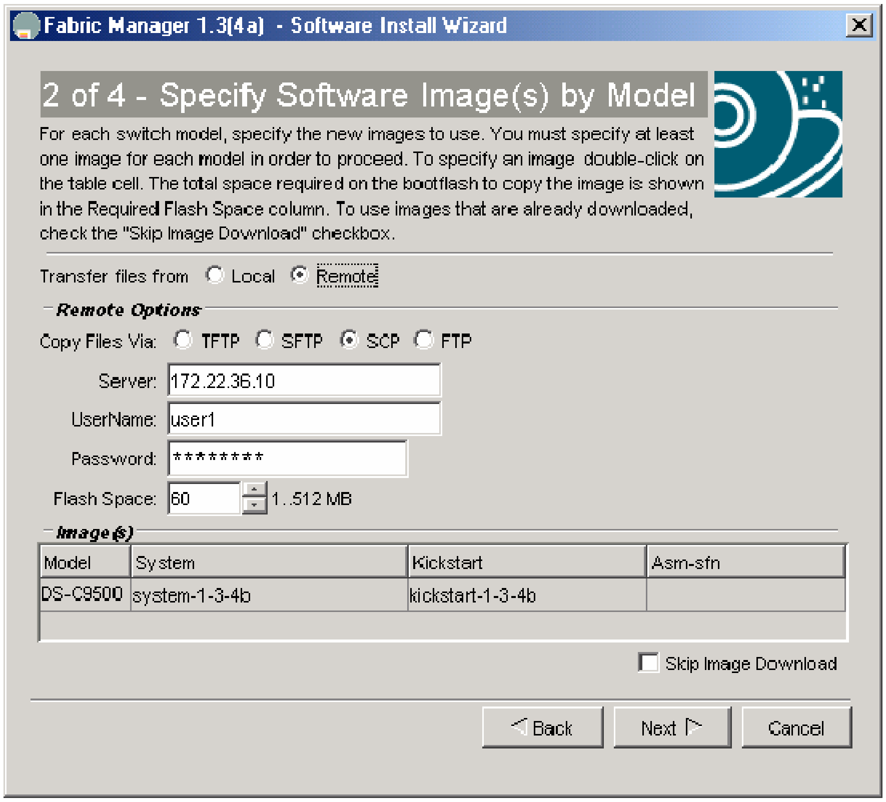

Fabric Manager 1.3(4a) - Software Install Wizard

2 of 4 - Specify Software Image(s) by Model

For each switch model, specify the new images to use. You must specify at least one image for each model in order to proceed. To specify an image double-click on the table cell. The total space required on the bootflash to copy the image is shown in the Required Flash Space column. To use images that are already downloaded, check the "Skip Image Download" checkbox.

Transfer files from ○ Local ● Remote

Remote Options

Copy Files Via: ○ TFTP ○ SFTP ● SCP ○ FTP

Server: 172.22.36.10

UserName: user1

Password: ********

Flash Space: 60 1..512 MB

Image(s)

Model	System	Kickstart	Asm-sfn
DS-C9500	system-1-3-4b	kickstart-1-3-4b	

□ Skip Image Download

Step 4 Click **Next**.

Depending on the installation method (that is, already downloaded to bootflash or download during the install), the wizard may prompt for additional file locations. The fourth and final screen provides a summary and enables you to start the install. During the installation, a compatibility screen pops up and displays the same version compatibility information that was displayed in the CLI upgrade. You must click **Yes** to continue with the upgrade.

Note Unlike the CLI, the Fabric Manager maintains connectivity to the switch and provides detailed information during the entire upgrade sequence, without requiring you to manually reestablish connectivity to the switch during the supervisor switchover. If there is a failure, the last screen displays any reasons for a failed upgrade.

Recovering a Password

If there are no accounts accessible on the Cisco MDS 9000 switch that have either network-admin or user account creation privileges, you may have to perform a password recovery on the admin account if passwords are lost.

Warning **This procedure requires console access to the switch and requires a reboot of the switch.**

Tip It is possible for another CLI user with network-admin privileges to change the password of the admin user, which can alleviate reloading the switch.

To recover the admin account's password, follow these steps:

Step 1 If possible, save the current configuration by entering the **copy-running config** command on the switch:

```
switch# copy running-config startup-config
[########################################] 100%
```

Step 2 Connect a console cable to the active supervisor of the MDS 9000 switch. (See Figure 2-4 and Figure 2-5.)

Figure 2-4 Console Connection on an MDS 9500 Series Switch

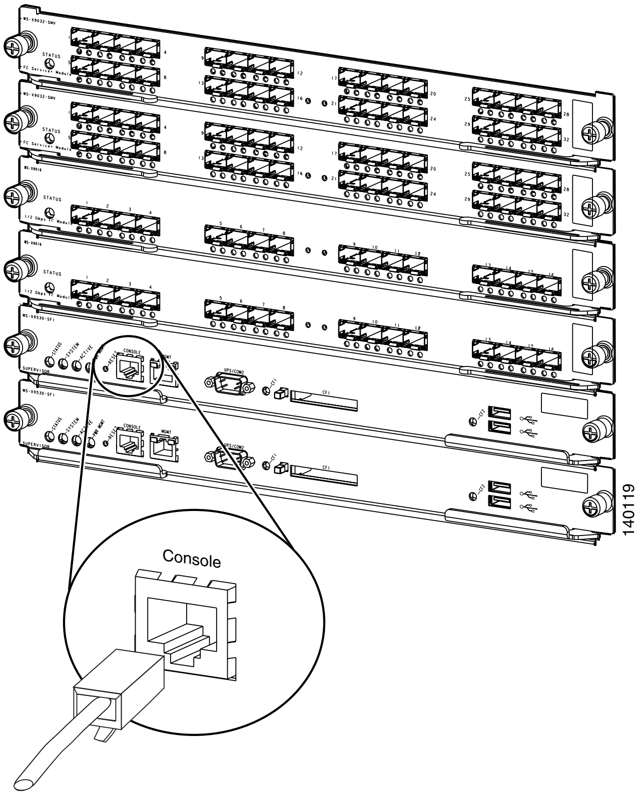

Figure 2-5 Console Connection on an MDS 9200 Series Switch

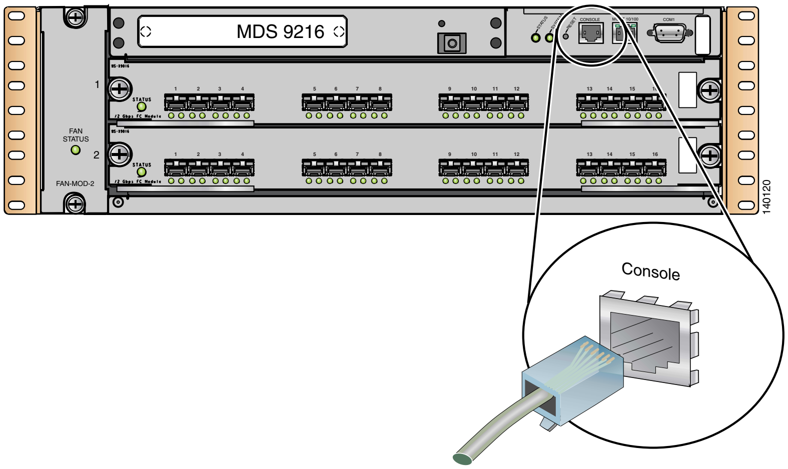

Step 3 Attach the RS-232 end of the console cable to a PC.

Step 4 Configure Hyperterm or a similar terminal emulation software for 9600 baud, 8 data bits, no parity, 1 stop bit and no flow control. (See Figure 2-6.)

Figure 2-6 HyperTerm Terminal Settings

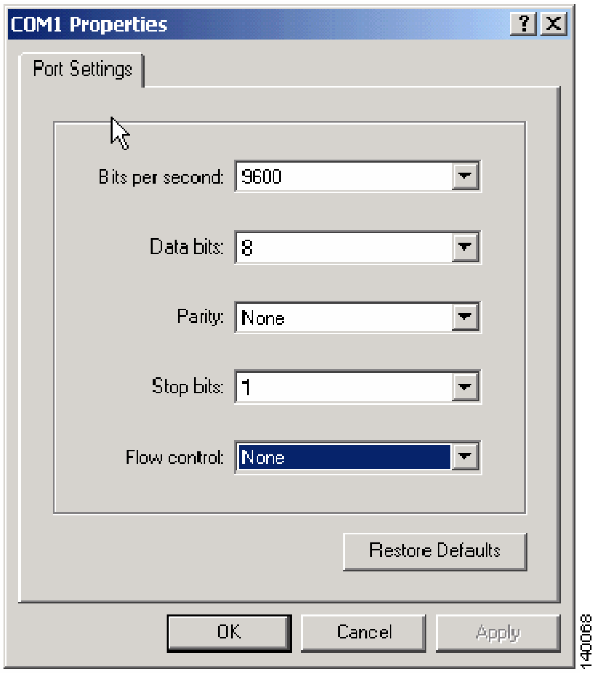
COM1 Properties

Port Settings

Bits per second: 9600

Data bits: 8

Parity: None

Stop bits: 1

Flow control: None

Step 5 Establish a connection to the switch if possible, at least enough to display the login prompt if no user accounts are available.

Step 6 For a multi-supervisor switch, MDS-9509 or MDS-9506, physically remove the standby supervisor. It is not necessary to remove it from the chassis, just enough so that it does not make contact with the backplane.

Step 7 Reboot the switch either by cycling the power or entering the **reload** command.

Step 8 Press **Ctrl-]** (when the switch begins its SAN-OS software boot sequence) to enter the `switch(boot)#` prompt.

Step 9 Enter configuration mode:

```
switchboot# config terminal
```

Step 10 Enter the **admin-password <new password>** command.

```
switch(boot-config)# admin-password temppassword
switch(boot-config)# exit
```

Step 11 Load the system image to finish the boot sequence.

```
switch(boot)# load bootflash: m9500-sf1ek9-mz.1.3.4b.bin
```

Step 12 Log in to the switch using the admin account and the temporary password.

```
switch login: admin
Password:
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2004, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or 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#
```

Step 13 Change the admin password to a new permanent password.

```
ca-9506# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
ca-9506(config)# username admin password g05ox
```

Step 14 Save the configuration that includes the new password.

```
switch# copy running-config startup-config
[########################################] 100%
```

Installing a License

To install a license key, use either the CLI and or the Fabric Manager.

Using the CLI to Install a License

Step 1 Copy the license file to the bootflash of the supervisor.

```
switch# copy scp://user1@172.22.36.10/tmp/FM_Server.lic bootflash:FM_Server.lic
user1@172.22.36.10's password:
FM_Server.lic   100% |*****************************|  2035    00:00
```

Step 2 Verify the license file.

```
switch# show license file FM_Server.lic
lic.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT FM_SERVER_PKG cisco 1.0 permanent uncounted \
       VENDOR_STRING=MDS HOSTID=VDH=FOX0713037X \
       NOTICE="<LicFileID>lic_template</LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=D8CF07EA26C2
```

Step 3 Cross reference the switch’s host-id (VDH=FOX0713037X) with that listed in the license file.

```
ca-9506# show license host-id
License hostid: VDH=FOX0713037X
```

Step 4 Install the license file.

```
switch# install license bootflash:FM_Server.lic
Installing license ..done
```

Step 5 Verify the license has been installed.

```
switch# show license
lic.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT FM_SERVER_PKG cisco 1.0 permanent uncounted \
       VENDOR_STRING=MDS HOSTID=VDH=FOX0713037X \
       NOTICE="<LicFileID>lic_template</LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=D8CF07EA26C2
```

Step 6 Display a summary of the installed licenses by issuing the **show license usage** command.

```
switch# show license usage
Feature               Insta License Status Expiry Date Comments
                      lled  Count
-----------------------------------------------------------------
FM_SERVER_PKG          Yes      -   In use never       -
MAINFRAME_PKG          No       -   Unused             -
ENTERPRISE_PKG         Yes      -   In use never       -
SAN_EXTN_OVER_IP       Yes      2   In use never       -
SAN_EXTN_OVER_IP_IPS4  No       0   Unused             -
-----------------------------------------------------------------
```

Step 7 Display the features within a license package are being used by specifying the package name. In this case QoS is using the Enterprise package.

```
ca-9506# show license usage ENTERPRISE_PKG
Application
-----------
Qos Manager
-----------
```

Using Fabric Manager to Install a License

To install a license, follow these steps:

Step 1 Click the **License Install** icon in the main toolbar of Fabric Manager to launch the License Installation Wizard. You see the License Install Wizard dialog box. (See Figure 2-7.)

Figure 2-7 License Installation Wizard

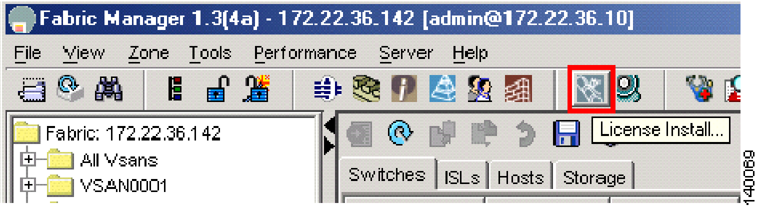

Step 2 In the License Install Wizard dialog box, as shown in Figure 2-8, check the appropriate check box to specify how to install the keys based upon whether or not you have already obtained the license key files or if you have only a Product Authorization Key (PAK). If you have a PAK, then you can download the license file and install it from the Cisco website.

Figure 2-8 License Installation Method

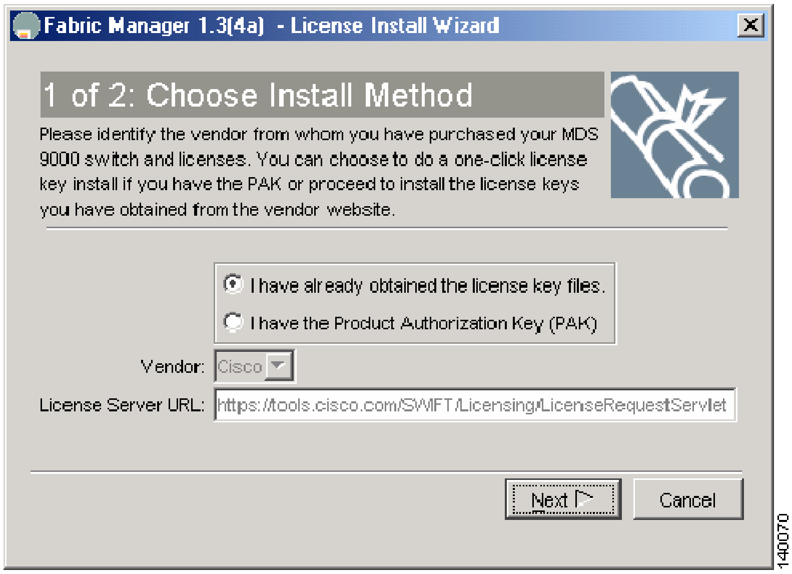

Fabric Manager 1.3(4a) - License Install Wizard

1 of 2: Choose Install Method

Please identify the vendor from whom you have purchased your MDS 9000 switch and licenses. You can choose to do a one-click license key install if you have the PAK or proceed to install the license keys you have obtained from the vendor website.

● I have already obtained the license key files.
○ I have the Product Authorization Key (PAK)

Vendor: Cisco

License Server URL: https://tools.cisco.com/SWIFT/Licensing/LicenseRequestServlet

Step 3 If the keys already exist on a server, enter the name and location of the license key files in the dialog box like the one in Figure 2-9.

Figure 2-9 License File Location

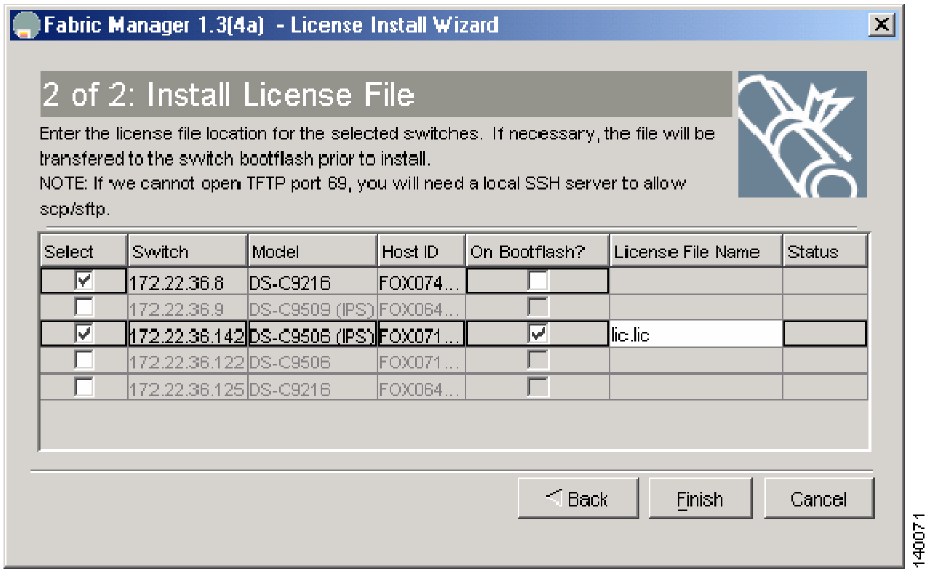

2 of 2: Install License File

Enter the license file location for the selected switches. If necessary, the file will be transfered to the switch bootflash prior to install.
NOTE: If we cannot open TFTP port 69, you will need a local SSH server to allow scp/sftp.

Select	Switch	Model	Host ID	On Bootflash?	License File Name	Status
☑	172.22.36.8	DS-C9216	FOX074...	☐		
☐	172.22.36.9	DS-C9509 (IPS)	FOX064...	☐		
☑	172.22.36.142	DS-C9506 (IPS)	FOX071...	☑	lic.lic	
☐	172.22.36.122	DS-C9506	FOX071...	☐		
☐	172.22.36.125	DS-C9216	FOX064...	☐		

If the license files are not already available, and you only have the PAK numbers, then Fabric Manager can obtain the license files directly from Cisco.com. (See Figure 2-10.)

Figure 2-10 Install License Using PAK

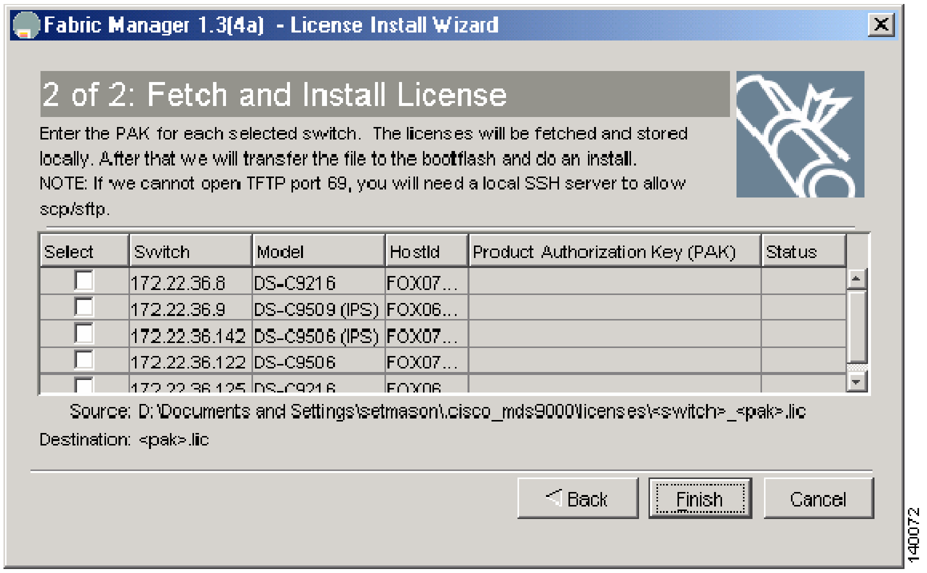

2 of 2: Fetch and Install License

Enter the PAK for each selected switch. The licenses will be fetched and stored locally. After that we will transfer the file to the bootflash and do an install.
NOTE: If we cannot open TFTP port 69, you will need a local SSH server to allow scp/sftp.

Select	Switch	Model	HostId	Product Authorization Key (PAK)	Status
☐	172.22.36.8	DS-C9216	FOX07...		
☐	172.22.36.9	DS-C9509 (IPS)	FOX06...		
☐	172.22.36.142	DS-C9506 (IPS)	FOX07...		
☐	172.22.36.122	DS-C9506	FOX07...		
☐	172.22.36.125	DS-C9216	FOX06		

Source: D:\Documents and Settings\setmason\.cisco_mds9000\licenses\<switch>_<pak>.lic
Destination: <pak>.lic

At this point, the license keys can be installed and the licensable feature can be used.

Copying Core Files from the MDS 9000 Switch

If an MDS 9000 switch process crashes, it may create a core file which you can send to Cisco TAC for further troubleshooting. To copy a core file off of the MDS 9000 switch, follow these steps:

Step 1 Before copying a core file to another server, identify the PID of the core file:

```
switch# show cores
Module-num Process-name PID Core-create-time
---------- ------------ --- ----------------
5            fspf       1524    Jul 15 03:11
```

Step 2 Copy the core file using FTP, for example, with the following command syntax:

```
"core://<module-number>/<process-id>"

switch# copy core://5/1524 ftp://172.22.36.10/tmp/fspfcore
```

You can now send the file to Cisco TAC according to the directions you receive from a TAC engineer.

Configuring an NTP Server

Network Time Protocol (NTP) is a protocol used by devices to synchronize their internal clocks with other devices. The Cisco MDS 9000 switch can only be used as an NTP client and can talk to other NTP systems which are considered to have a higher stratum (or authority). NTP is hierarchical in nature such that the lower stratum numbers are closer to the source of the time authority. Devices that are at the same stratum can be configured as peers so that they can work together to determine the correct ime by making minute adjustments. Normally, the MDS 9000 switches are configured as peers, while a router or other dedicated machine is used as an NTP server.

Note NTP will not set the time zone (or offset from UTP) for the switch. You must manually set the time zone using, for example, Eastern Standard Time and Eastern Daylight-Savings Time:
clock timezone EST -5.0
clock summer-time EDT 1 Sunday Apr 02:00 5 Sunday Oct 02:00 60

The following example uses these IP addresses:

Switch #1 IP Address: 172.22.36.142

Switch #2 IP Address: 172.22.36.9

NTP Server: 171.69.16.26

To configure NTP for switch1, follow these steps:

Step 1 Enter configuration mode and add the NTP server.

```
switch1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# ntp server 171.69.16.26
```

Step 2 Add the NTP peer switch.

```
switch1(config)# ntp peer 172.22.36.9
switch1(config)# end
```

At this point, NTP is configured and the switch will slowly adjust to the new time.

Step 3 To view the NTP configuration, enter the **show ntp peers** command:

```
switch1# show ntp peers
--------------------------------------------------
  Peer IP Address               Serv/Peer
--------------------------------------------------
  171.69.16.26                  Server
  172.22.36.9                   Peer
```

Restoring a Fixed Switch Configuration

This procedure covers the process of backing up and restoring a switch configuration for one of the Cisco MDS 9000 Family switches that have a fixed configuration. These include the Cisco MDS 9216 and 9100 series fabric switches.

This procedure leverages the following resources:

- Old Switch: switch1: (172.22.36.8)
- New Switch: switch2
- File Server: host1

Note Only restore a switch configuration to a switch that has the exact same firmware version on it as was used to create the switch configuration. If an upgrade is required, restore the configuration, and then upgrade the firmware.

To restore a fixed switch configuration, follow these steps:

Step 1 Save the running configuration using the following command.

```
switch1# copy running-config startup-config
[########################################] 100%
```

Step 2 Copy the startup configuration to the file server using any of the available methods on the MDS 9000 switch, such as FTP, TFTP, SFTP, or SCP.

```
switch1# copy startup-config scp://user@host1/switch1.config
user@switch1's password:
sysmgr_system.cfg    100% |*****************************| 10938       00:00
switch1#
```

Step 3 Capture the port assignments using the fabric login (FLOGI) database. The database is used to verify that all of the cables are placed in the correct locations.

```
switch1# show flogi database
---------------------------------------------------------------------------
INTERFACE  VSAN    FCID            PORT NAME               NODE NAME
---------------------------------------------------------------------------
fc1/8      600   0x7c0007  50:05:07:63:00:ce:a2:27  50:05:07:63:00:c0:a2:27
fc1/13     1001  0xef0001  50:06:0e:80:03:4e:95:13  50:06:0e:80:03:4e:95:13
fc1/15     600   0x7c0004  50:06:0b:00:00:13:37:ae  50:06:0b:00:00:13:37:af
```

Note At this point, the old switch is no longer needed; its mgmt0 port should be disconnected from the LAN.

Step 4 Log on to the new switch using the console connection and clear the switch configuration. Do not run the setup script, if prompted. The **write erase** command will erase the switch configuration.

```
switch2# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n)  [n] y
```

Step 5 Reload the switch.

```
switch2# reload
This command will reboot the system. (y/n)?  [n] y
```

When the switch comes up in its factory default mode and prompts for the Basic System Configuration Dialog, skip it because all the configuration options are contained in the startup configuration file of the old switch.

Step 6 Manually configure the IP address as follows.

```
switch2# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch2(config)# int mgmt 0
switch2(config-if)# ip address 172.22.36.8 255.255.254.0
switch2(config-if)# no shut
```

Step 7 If interface (fc X/Y) based zoning is used, obtain the WWN for the new switch. Otherwise, skip this step.

```
switch2# show wwn switch
Switch WWN is 20:00:00:0d:ec:02:1d:40
```

Step 8 On the file server, make a copy of the configuration file and then open it in a text editor, such as Notepad or vi.

a. Remove the lines that contain the SNMP user accounts, as the encrypted passwords are tied to the MAC address of the chassis:

```
$ cp switch1.config switch1.config.orig
$ vi switch1.config
```

The user accounts are all grouped together and begin with **snmp-server user**:

```
snmp-server user admin network-admin auth md5 0x46694cac2585d39d3bc00c8a4c7d48a6
 localizedkey
snmp-server user guestadmin network-admin auth md5 0xcae40d254218747bc57ee1df348
26b51 localizedkey
```

b. If interface (fc X/Y) zoning was not used, skip this step. Otherwise, replace the WWN of the old switch in the zone member commands with the WWN of the new switch:

```
zone name Z_1 vsan 9
  member interface fc1/9 swwn 20:00:00:0d:ec:02:1d:40
```

c. Save and exit the configuration file.

Step 9 From the new switch, copy the modified configuration file from the file server onto the running configuration of the new switch. As the file is copied, it executes on the switch as the configuration is applied. The commands being applied are contained in single quotes. Any errors caused by applying the commands are displayed immediately after the error-causing command executes. The prompt changes to reflect the new switch name.

```
switch2# copy scp://user@host1/switch1.config running-config
user@host1's password:
switch1.config   100% |*****************************| 10938       00:00
```

Step 10 Save the configuration by copying startup-config to running-config.

```
switch1# copy running-config startup-config
[########################################] 100%:
```

Step 11 At this point, access the switch via the CLI and perform the following remaining items:

a. Recreate SNMP user accounts.

b. Remove the MDS 9000 switch entry from the host's known_hosts file, because the switch's public key is different.

c. Install license keys, if required.

Step 12 Move the cables from the old switch to the new switch, using the **show flogi database** command on the old switch as a reference to verify that each cable is in the correct location.

Step 13 Verify that all devices have logged in and that all features are running as they are supposed to be.

Step 14 Save the running-configuration to the startup-configuration with the **copy running-config startup-config** command.

Step 15 Reload the switch to verify that it boots correctly with the configuration.

Preparing to Call Cisco TAC

At some point, the administrator may need to contact the Cisco TAC or their OSM for some additional assistance. This section outlines the steps that the administrator should perform prior to contacting their next level of support, as this will reduce the amount of time needed to resolve the issue.

Step 1 Do not reload the line card or the switch until you have completed at least Step 2. Some logs and counters are kept in volatile storage and will not survive a reload.

Step 2 Collect switch information and configuration. Do this before the issue is resolved and after it is resolved. The following three methods of collecting switch information each provide the same information.

a. CLI: Configure the Telnet/SSH application to log the screen output to a text file and issue the **show tech-support details** command.

b. CLI: Issue the **tac-pac <filename>** command, as in this example: `tac-pac bootflash://showtech.switch1`.

The **tac-pac** command redirects the output of a show **tech-support details** command to a file that you can then gzip. If no filename is specified, the file created is volatile:show_tech_out.gz. Copy the file off the MDS 9000 switch using the procedure described in Copying Files to or from the MDS 9000 Switch, page 2-2.

c. Fabric Manager: Choose **Tools** > **Show tech support**. Fabric manager can capture switch configuration information from multiple switches simultaneously. The file can be saved on the local PC.

Step 3 Capture the exact error codes:

a. If the error occurs in Fabric Manager, take a screen shot of the error. In Windows, use **ALT+Print Screen** to capture the active window, or press the **Print Screen** key to capture the entire desktop. Paste the screen capture into a new MSpaint.exe (or similar program) session.

b. Display the message log using the **show logging log** command or view the last X lines of the log using the **show logging last** *lines* command.

Step 4 Answer the following questions before placing a call to TAC:

a. In which switch, HBA, or storage port is the problem occurring? List MDS firmware, driver versions, operating systems versions and storage device firmware.

b. What is the network topology? (In Fabric Manager, choose **Tools -> Show Tech**. Save the map.)

c. Were any changes being made to the environment (zoning, adding line cards, upgrades) prior to or at the time of this event?

d. Are there other similarly configured devices that could have this problem but do not have it?

e. Where is this problematic device connected (MDS 9000 switch Z, interface x/y)?

f. When did this problem first occur?

g. When did this problem last occur?

h. How often does this problem occur?

i. How many devices have this problem?

j. Were any traces or debug outputs captured during the problem time? What troubleshooting steps have already been done? Were any of the following tools used?

- Fcanalyzer, PAA-2, Ethereal, local or remote SPAN
- CLI debug commands
- FC traceroute, FC ping
- FM/DM

Implementing Syslog

The syslog message server allows Cisco MDS 9000 switches to send a copy of the message log to a host for more permanent storage. Saving the logs in this way can be useful if the logs need to be examined over a long period of time or when the MDS 9000 switch is not accessible.

This example demonstrates how to configure a Cisco MDS 9000 switch to use the syslog facility on a Solaris platform. Although a Solaris host is being used, the syslog configuration on all UNIX and Linux systems is very similar.

Syslog uses the concept of a facility to determine how a message 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 sent via e-mail to a particular user. Specifying a severity determines that all messages of that level and greater severity (lower number) will be acted upon.

Tip The MDS 9000 switch messages should be logged to a different file from the standard syslog file so that they cannot be confused with other non-MDS 9000 switch syslog messages. To prevent log messages from filling up the /filesystem directory, do not locate the log file on the /filesystem directory.

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

To configure a Cisco MDS 9000 switch to use the syslog facility on a Solaris platform, follow these steps:

Step 1 Configure the MDS 9000 switch using the **config terminal** command:

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

Step 2 Display the configuration using the **show logging server** command:

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

Step 3 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 syslogd:

```
# /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 4 Test the syslog server by creating an event on the MDS 9000 switch. In this case, port fc1/2 was bounced and the following information 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)
```

CHAPTER 3

Physical Interfaces

This chapter describes options you can use to configure Fibre Channel (FC) and Gigabit Ethernet ports on an MDS multi-protocol switch. It includes the following sections:

- Configuring Fibre Channel Ports, page 3-1
- Configuring Gigabit Ethernet Ports, page 3-5

Configuring Fibre Channel Ports

The recipes in this section show how to configure various parameters and modes for a physical port on a Cisco MDS 9000 switch.

Port Description

The port description allows you to provide a plain text name as a description for the interface. In the following example, the Fibre Channel interface fc 1/1 is given the description “storage array 17 port 1,” which then becomes the port name on the switch.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport description “storage array 17 port 1”
switch(config-if)# end
switch#
```

Port Speed

This example shows how to the set the port speed for fc 1/1 to 1 GB, 2 Gb or auto-negotiate speed.

Note A port can be set to one speed. The default is auto-negotiate.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport speed 1000 <- port speed set to 1 GB
switch(config-if)# switchport speed 2000 <- port speed set to 2 GB
switch(config-if)# switchport speed auto <- port speed set to auto negotiate
switch(config-if)# ^Z
switch#
```

Port Mode (auto)

Note A Fibre Channel port can be set to one port mode at any given time. The default is auto on the 16-port line cards and FX on the 32-port line cards.

Setting the port mode to auto allows the port to negotiate to either an F, FL, or E port. It *cannot* negotiate to either ST, SD, or TL port modes. Auto port setting is the default setting for all ports on a 16-port line card. The following example shows how to set the port fc 1/1 to auto port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode auto
switch(config-if)# end
switch#
```

Port Mode (E)

Setting the port mode to E restricts the port to coming up as an E port (either trunking or non-trunking, depending on the trunking port mode). E port mode is used when the port functions as one end of an Inter-Switch Link (ISL). The following example shows how to set the port fc 1/1 to E port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode E
switch(config-if)# end
```

Port Mode (F)

Setting the port mode to F restricts the port to coming up as an F port. F port mode is used for end devices which can only communicate in point-to-point mode or to a switch. The following example shows how to set the port fc 1/1 to F port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode F
switch(config-if)# end
switch#
```

Port Mode (FL)

Setting the port mode to FL restricts the port to coming up as an FL port. FL port mode is used for end devices which can only communicate as a public loop device. The following example shows how to set the port fc 1/1 to FL port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode FL
switch(config-if)# end
switch#
```

Port Mode (Fx)

Setting the port mode to Fx restricts the port to coming up as an F or FL port. Fx port mode is used exclusively for end devices and prevents a port from autonegotiating to an E port. The following example shows how to set the port fc 1/1 to Fx port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode Fx
switch(config-if)# end
switch#
```

Port Mode (SD)

Setting the port mode to SD configures the port for usage with a span session as the span destination (SD) port. This setting is used in conjunction with the PAA to span a port and obtain Fibre Channel traces without a Fibre Channel analyzer. The following example shows how to set the port fc 1/1 to SD port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode SD
switch(config-if)# end
```

Port Mode (ST)

Setting the port mode to ST configures the port for usage with a remote span session as the span tunnel (ST) port. This setting is used to set up a remote SPAN session to a remote switch in which a PAA or protocol analyzer is connected. The following example shows how to set the port fc 1/1 to ST port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode ST
switch(config-if)# end
switch#
```

Port Mode (TL)

Setting the port mode to TL restricts the port to coming up as a TL port. TL port mode is used exclusively for end devices which can only communicate as a private loop device. The following example shows how to set the port fc 1/1 to TL port mode:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport mode TL
switch(config-if)# end
switch#
```

Configuring Trunking E ports

A trunking port carries VSAN enabled frames between switches. The section below shows the various configuration options for a trunking port.

Note These same principles apply to PortChannels. Just specify the PortChannel interface, **int port-channel 1**, rather than an individual link, **interface fc 1/1**.

Setting the Trunk Port Mode

This example shows how to the set the trunk port mode for fc 1/1 to auto, on and off. The default mode is auto. One end of an ISL should be set to on when connected between two MDS switches, while the other end can be either on or auto. The trunk port mode needs to be off when the ISL is talking to non-MDS switches.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport trunk mode auto <- auto negotiates trunk port mode
switch(config-if)# switchport trunk mode on   <- sets trunk port mode to on
switch(config-if)# switchport trunk mode off  <- sets trunk port mode to off
switch(config-if)# ^Z
switch#
```

Configuring Trunk Ports to Filter Specific VSANs

This example shows how to configure the allowed VSAN traffic through the interface fc 1/1. The **all** keyword allows all VSAN traffic to go through the port. The **add 2** keyword adds VSAN 2 to the list of allowed VSANs through the port. The **add 2-4** keyword adds VSANs 2 to 4 to the list of allowed VSANs through the port. The default mode is to allow all the VSAN traffic to pass through the port.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport trunk allowed vsan all     <- all VSAN traffic
switch(config-if)# switchport trunk allowed vsan add 2   <- only VSAN 2 traffic
switch(config-if)# switchport trunk allowed vsan add 2-4 <- VSAN 2 to 4 traffic
switch(config-if)# ^Z
switch#
```

Enabling Port Beaconing

This example shows how to enable the LEDs below the port fc 1/1 to start flashing. The flashing is useful in identifying a port for physical cabling or troubleshooting.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc 1/1
switch(config-if)# switchport beacon
switch(config-if)# end
```

Configuring Gigabit Ethernet Ports

Configuring VRRP

Virtual Router Redundancy Protocol (VRRP) allows two Gigabit Ethernet interfaces to provide failover capability for an IP address. The two interfaces will form an active/passive or master/backup state in which one interface will service requests for the shared IP address while the other will remain in a backup or standby state. VRRP is ideal for providing port level redundancy in iSCSI configurations. A Gigabit Ethernet port can still have its own IP address while partaking in a VRRP configuration.

A VRRP session has an ID assigned to it. Using this ID, the two interfaces will communicate to identify its peer. The same ID must be used on both switches. The procedure for having both members of the VRRP pair on the same switch would be the same as if the two members were on different switches.

Note To have one interface become the master interface whenever it is online, which is known as preemption, set the Gigabit Ethernet interface to have the same IP address as the VRRP IP address.

In this example, the configuration is as follows:

VRRP ID: 1

VRRP IP address: 192.168.1.40

Switch 1:

Interface gige3/3 (192.168.1.20)

Switch 2:

Interface gige4/1 (192.168.1.30)

To configure VRRP, follow these steps:

Step 1 Configure the IP address on the two Gigabit Ethernet interfaces.

```
Switch1# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch1(config)# interface gigabitethernet 3/3
Switch1(config-if)# ip address 192.168.1.20 255.255.255.0
Switch1(config-if)# no shut

Switch2# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch2(config)# interface gigabitethernet 4/1
Switch2(config-if)# ip address 192.168.1.30 255.255.255.0
Switch2(config-if)# no shut
```

At this point, it is a good idea to verify that a host on the local subnet can ping both IP addresses (192.168.1.20 and 192.168.1.30). Alternatively, use the **ips measure-rtt** command to ping from one Gigabit Ethernet port to the other.

```
Switch1# ips measure-rtt 192.168.1.30 interface gigabitethernet 3/3
Round trip time is 172 micro seconds (0.17 milli seconds)
```

Step 2 Configure the VRRP session on both switches using the VRRP ID (1).

```
Switch1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Switch1(config)# interface gigabitethernet 3/3
Switch1(config-if)# vrrp 1
Switch1(config-if-vrrp)# address 192.168.1.40
Switch1(config-if-vrrp)# no shut
Switch1(config-if-vrrp)# end

Switch2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Switch2(config)# interface gigabitethernet 4/1
Switch2(config-if)# vrrp 1
Switch2(config-if-vrrp)# address 192.168.1.40
Switch2(config-if-vrrp)# no shut
Switch2(config-if-vrrp)# end
```

Step 3 Verify that the VRRP session is up and which interface has become the master. Enter the following commands:

```
Switch2# show vrrp vr 1
                       Interface    VR     Status
-------------------------------------------------------
              GigabitEthernet3/3     1     backup

switch1# show vrrp vr 1 interface gig3/3 status
vr id 1 status
MAC address 00:00:5e:00:01:01
Operational state: master
Up time 8 sec
```

To view the configuration, enter the following command:

```
Switch1# show vrrp vr 1 interface gigabitethernet 3/3 configuration
vr id 1 configuration
admin state up
priority 100
associated ip: 192.168.1.40
no authentication
advertisement-interval 1
preempt no
protocol IP
```

CHAPTER 4

Logical Interfaces

PortChannels allow for the aggregation of multiple Fibre Channels or FCIP links into a single, high-speed, fault tolerant Fibre Channel or FCIP ISL (Inter-Switch Link). An ISL with multiple Fibre Channels or FCIP links has the same configuration options as a single-link Fibre Channel or FCIP ISL. This chapter focuses on the procedures associated with building, modifying, and reducing a PortChannel. It includes the following sections:

Creating a PortChannel

There are two basic methods to create a PortChannel: using the Fabric Manager wizard or entering commands through the CLI (command-line interface). This procedure shows how to create a PortChannel using the CLI.

This procedure uses the following resources:

Switch 1:

Channel Group 2

Interfaces: fc1/1 and fc2/1

Switch 2:

Channel Group 2

Interfaces: fc1/1 and fc2/1

Allowed VSANs: 1,10

Figure 4-1 Topology

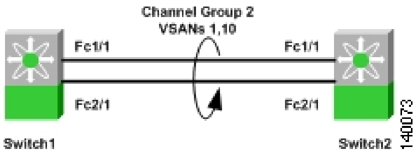

Tip

- A PortChannel should use interfaces on multiple line cards to protect the PortChannel against line card failure.
- The same channel group number should be used on both ends of a PortChannel. Having the same number helps in troubleshooting and identifying the corresponding channel group on the other switch.
- A PortChannel, like all other interfaces, can have a description. Use the description field to specify exactly where the PortChannel goes.
- PortChannels can use any port on the switch and connect to any other port on a switch.
- Set the initial VSAN allowed list prior to bringing up the PortChannel. This action prevents any VSANs from merging when you bring up the PortChannel the first time.

To create a PortChannel through the CLI, follow these steps:

Step 1 Create the PortChannel on switch 1.

```
switch1# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# interface fc1/1, fc2/1
switch1(config-if)# channel-group 2
fc1/1 fc2/1 added to port-channel 2 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring them up
switch1(config-if)# switchport description "To switch2 PortChannel2"
```

Step 2 Enable trunking (TE) and set the VSAN allowed list on switch 1.

```
switch1# config terminal
switch1(config)# int port-channel 2
switch1(config-if)# switchport trunk mode on
switch1(config-if)# switchport trunk allowed vsan 1
switch1(config-if)# switchport trunk allowed vsan add 10
```

Step 3 Create the PortChannel and set the description on switch 2.

```
switch2# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch2(config)# interface fc1/1, fc2/1
switch2(config-if)# channel-group 2
fc1/1 fc2/1 added to port-channel 2 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring them up
switch2(config-if)# switchport description "To switch1 PortChannel2"
```

Step 4 Enable trunking (TE) and set the VSAN allowed list on switch 2.

```
switch2# config terminal
switch2(config)# int port-channel 2
switch2(config-if)# switchport trunk mode on
switch2(config-if)# switchport trunk allowed vsan 1
switch2(config-if)# switchport trunk allowed vsan add 10
```

Step 5 Enable the interfaces to bring up the PortChannel.

a. Switch 1

```
switch1# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# interface fc1/1, fc2/1
switch1(config-if)# no shut
```

b. Switch 2

```
switch2# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch2(config)# interface fc1/1, fc2/1
switch2(config-if)# no shut
```

Step 6 Verify that the PortChannel came up.

```
switch1# show interface port-channel 2
port-channel 2 is trunking
    Port description is To switch2 PortChannel2
    Hardware is Fibre Channel
    Port WWN is 24:02:00:0c:85:e9:d2:c0
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 4 Gbps
    Trunk vsans (admin allowed and active) (1,10)
    Trunk vsans (up)                       (1,10)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
    5 minutes output rate 56 bits/sec, 7 bytes/sec, 0 frames/sec
      78296342 frames input, 72311141128 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      56299070 frames output, 26061293700 bytes
        0 discards, 0 errors
      0 input OLS, 2 LRR, 0 NOS, 0 loop inits
      4 output OLS, 2 LRR, 0 NOS, 0 loop inits
    Member[1] : fc1/1
    Member[2] : fc2/1
    iSCSI authentication: None
```

Adding a New Member to a PortChannel

This procedure uses the following resources:

Switch 1:

Channel Group 2

Existing Interfaces: fc1/1 and fc2/1

New Interfaces: fc3/1

Switch 2:

Channel Group 2

Existing Interfaces: fc1/1 and fc2/1

New Interface: fc3/1

To add a member to an existing PortChannel, follow these steps:

Step 1 To add the new member to switch 1, use the **force** keyword to have the new link inherit the parameters of the existing link.

```
switch1# conf t
switch1(config)# int fc3/1
switch1(config-if)# channel-group 2 force
fc3/1 added to port-channel 2 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring them up
switch1(config-if)# no shut
```

Step 2 To add the new member to switch 2, use the **force** keyword to have the new link inherit the parameters of the existing link.

```
switch2# conf t
switch2(config)# int fc3/1
switch2(config-if)# channel-group 2 force
fc3/1 added to port-channel 2 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring them up
switch3(config-if)# no shut
```

Step 3 Verify that the PortChannel now has all three members.

```
switch1# show interface port-channel 2
port-channel 2 is trunking
    Port description is To switch2 PortChannel2
    Hardware is Fibre Channel
    Port WWN is 24:02:00:0c:85:e9:d2:c0
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 6 Gbps
```

```
Trunk vsans (admin allowed and active) (1,10)
Trunk vsans (up)                       (1,10)
Trunk vsans (isolated)                 ()
Trunk vsans (initializing)             ()
5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
5 minutes output rate 56 bits/sec, 7 bytes/sec, 0 frames/sec
  78296342 frames input, 72311141128 bytes
    0 discards, 0 errors
    0 CRC,  0 unknown class
    0 too long, 0 too short
  56299070 frames output, 26061293700 bytes
    0 discards, 0 errors
  0 input OLS, 2 LRR, 0 NOS, 0 loop inits
  4 output OLS, 2 LRR, 0 NOS, 0 loop inits
Member[1] : fc1/1
Member[2] : fc2/1
Member[3] : fc3/1
iSCSI authentication: None
```

Removing a Member from a PortChannel

This procedure shows how to remove a member from a PortChannel.

Tip Using the **quiesce** command allows seamless reduction in the PortChannel, with no dropping of in-flight frames.

This procedure uses the following resources:

Switch 1:

Channel Group 2

Existing Interfaces: fc1/1, fc2/1, fc3/1

Remove Interface: fc3/1

Switch 2:

Channel Group 2

Existing Interfaces: fc1/1, fc2/1, fc3/1

Remove Interface: fc3/1

To remove a member of a PortChannel, follow these steps:

Step 1 On switch 1, quiesce the interfaces to be removed.

```
switch1# quiesce interface fc3/1
WARNING: this command will stop forwarding frames to the specified interfaces. I
t is intended to be used to gracefully shutdown interfaces in a port-channel. The
procedure is:
1. quiesce the interfaces on both switches.
2. shutdown the interfaces administratively.
Do you want to continue? (y/n)  [n] y
fc3/1: quiesced
Please quiesce the corresponding interfaces on the other switch and then shut down
them administratively.
```

Step 2 On switch 2, quiesce the interfaces to be removed.

```
switch2# quiesce interface fc3/1
WARNING: this command will stop forwarding frames to the specified interfaces. I
t is intended to be used to gracefully shutdown interfaces in a port-channel. The
procedure is:
1. quiesce the interfaces on both switches.
2. shutdown the interfaces administratively.
Do you want to continue? (y/n)  [n] y
fc3/1: quiesced
Please quiesce the corresponding interfaces on the other switch and then shut down
them administratively.
```

Step 3 Shut the physical interfaces on both switches.

a. On Switch 1:

```
switch1# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# int fc3/1
switch1(config-if)# shut
```

b. On Switch 2:

```
switch2# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch2(config)# int fc3/1
switch2(config-if)# shut
```

Modifying the VSAN Allowed List

Modifying the VSAN allowed list for a PortChannel is the same as modifying it for a standard, single link TE port.

To add VSAN 17 to PortChannel 2, enter these commands:

```
switch2# config terminal
switch2(config)# int port-channel 2
switch2(config-if)# switchport trunk allowed vsan add 17
```

To remove VSAN 17 from PortChannel 2, enter these commands:

```
switch2# config terminal
switch2(config)# int port-channel 2
switch2(config-if)# no switchport trunk allowed vsan add 17
```

VSANs

This chapter provides recipes for configuring VSANs (virtual SANs). A VSAN is a logical grouping of ports in a single switch or across multiple switches that function like a single fabric.

This chapter includes the following sections:

VSAN Overview

A VSAN is a logical fabric. Each VSAN has all the required fabric services, independent of the other VSANs, configured on the same switch or set of switches. A VSAN provides:

- SAN island consolidation on a high-port-density physical switch
- Traffic isolation
- Increased security

VSANs can be numbered from 1 to 4094. VSAN 1 and VSAN 4094 are predefined and have very specific roles. VSAN 1 is the default VSAN which holds all the ports by default and the VSAN 4094 is the isolated VSAN into which orphaned ports are assigned.

Creating a VSAN on a Single Switch and Adding an Interface

This recipe shows the steps to create and name a VSAN on a single switch. Enter the following commands to create VSAN 200 with the name TapeVSAN and add Fibre Channel interface fc 1/1.

```
switch1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 name TapeVSAN
switch(config-vsan-db)# vsan 200 interface fc 1/1
switch(config-vsan-db)# ^Z
switch#
```

Setting VSAN Interop Mode

Interop mode can be set for VSANs that need to interact with other third-party switches. Interop mode 1 is required when all vendor switches are set in their respective interop modes. In interop mode, only domain IDs 97 to 127 are allowed. Interop mode 2 is required when a VSAN has to work with a Brocade 2800/3800 switch in its native mode. Interop mode 3 is required when the VSAN has to work with Brocade 3900 or a 12000 switch. For more information, refer to the *MDS Switch to Switch Interoperability Configuration Guide*.

Note Before performing any interoperability work, consult the *MDS Switch to Switch Interoperability Configuration Guide,* which explains and provides detailed instructions on using the interop modes. The guide may be found at the following location: http://www.cisco.com/en/US/products/hw/ps4159/ps4358/products_device_support_tables_list.html.

To set the interop modes 1, 2, 3 for a VSAN, follow these steps:

Step 1 For **Interop mode 1,** ensure that the domain ID of the VSAN is between 97 – 127 for this mode to work. To change the interop mode of VSAN 200 to interop mode 1, enter these commands:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 interop 1
switch(config-vsan-db)# ^Z
switch#
```

Step 2 For **Interop mode 2** with brocade 2800/ 3800 switches, change the interop mode of VSAN 200 to interop mode 2.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 interop 2
switch(config-vsan-db)# ^Z
switch#
```

Step 3 **For Interop mode 3** with brocade switches 3900/12000, change the interop mode of VSAN 200 to interop mode 3.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 interop 3
switch(config-vsan-db)# ^Z
switch#
```

Changing a Load Balancing Scheme

The load-balancing scheme can be configured per VSAN. On a Cisco MDS 9000 switch, you can configure S_ID (source id) or D_ID (destination id) based load balancing, as well as exchange level (S_ID, D_ID, OX_ID) load balancing.

Sequence Level load-balancing (Source_ID, Destination_ID)

To change the load-balancing scheme for a VSAN 200 to S_ID, D_ID mode, enter the following commands:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 loadbalancing src-dst-id
switch(config-vsan-db)# ^Z
switch#
```

Exchange Level Load Balancing (S_ID, D_ID, OX_ID)

To change the load-balancing scheme for a VSAN 200 to S_ID, D_ID, and OX_ID mode, enter the following commands:

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 loadbalancing src-dst-ox-id
switch(config-vsan-db)# ^Z
switch#
```

Exchange level load balancing is the default load-balancing scheme.

Configuring a Static Domain ID and Persistent FC ID

Within a VSAN, the domain manager process on the principal switch in a fabric is responsible for assigning a domain_ID to switch that is joining the fabric. When a switch boots up or joins a new fabric, it can request a specific domain_ID or take any available domain_ID. A domain_ID can be configured in one of three ways:

- **Dynamic** (default): The new switch will not request any domain_ID from the principal switch and accept any domain_ID that is assigned.
- **Preferred**: The new switch will request a specific domain_ID, however, if it receives a different domain_ID it will accept it.
- **Static**: The new switch will request a specific domain_ID, however, if it receives a different domain_ID, it will isolate itself from the fabric. This is the case when the same domain_ID must be maintained under all circumstances.

After obtaining the domain_ID from the principal switch in the VSAN, the local switch will assign Fibre Channel Identifiers (FC IDs) to each device as they log in to the Fabric. This process is known as FLOGI (Fabric Login). Some devices require that the same FC ID be assigned to a device as the FC ID used in the host's device path. HPUX and AIX are two operating systems that use the FC ID in the device path to the storage. To have the switch always assign the same FC ID to a device, a persistent FC ID must be configured for the VSAN. By default, the switch assigns the same FC ID to a device; however, if the switch is rebooted, the database of pWWN/FC ID mapping is not maintained. Enabling persistent FC IDs makes this database persistent.

When persistent FC ID is enabled, the MDS 9000 switch makes all of the devices in that VSAN persistent. Therefore, the admin is not required to manually type in devices entering that VSAN.

This recipe shows the steps to configure a static domain_ID for a VSAN and also how to enable a persistent FC ID for the same VSAN.

This procedure sets the domain_ID of VSAN 200 to 22 and then enables a persistent FC ID.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# fcdomain domain 22 static vsan 200
switch(config)# fcdomain fcid persistent vsan 200
switch(config)# ^Z
switch#
```

Note If the domain ID of VSAN 200 is different than what is currently running (22 in this case), then the VSAN will have to be restarted for the configuration changes to the domain_ID and FC ID persistence to take effect. Read the following Warning.

Warning **Changing domain_IDs and therefore FC IDs for a device is disruptive, because an end device will have to log in to the fabric (FLOGI) again to obtain a new FC ID. However, making a domain_ID static without changing its value is not disruptive.**

Restarting a VSAN

Sometimes the VSAN on a switch will need to be restarted. For example, after changing the domain_ID of a VSAN, the VSAN should be restarted for the new domain_ID to take effect.

The recipe shows how a VSAN (200) can be restarted.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vsan database
switch(config-vsan-db)# vsan 200 suspend
switch(config-vsan-db)# no vsan 200 suspend
switch(config-vsan-db)# end
```

Assigning a Predetermined FC ID to a PWWN

When performing a migration or HBA replacement, the same FC ID as was used previously may need to be assigned to the new pWWN. This recipe shows steps to assign a predetermined FC ID to a specific pWWN.

FC ID 0x160000 will be assigned permanently to pWWN 50:06:0b:82:bf:d1:db:cd. Therefore, when the pWWN logs into the switch (FLOGI), it will get this assigned FC ID.

Note The FC ID to be assigned (0x160000) should contain the same domain_ID (0x16) as the currently running domain in the VSAN.

```
switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan 22 wwn 50:06:0b:82:bf:d1:db:cd fcid 0x160000 dynamic
switch(config-fcid-db)# ^Z
switch#
```

CHAPTER 6

Inter-VSAN Routing

Inter-VSAN routing (IVR) facilitates the communication between a target and initiator in different VSANs, using IVR zones and IVR zone sets.

Cross VSAN communication is not permitted by default on a Cisco MDS 9000 switch. If there are two VSANs on a single switch, and one or more of the initiators in one VSAN needs to communicate with one or more of the targets in the other VSAN, IVR enables this communication.

This chapter describes how to configure IVR and includes the following sections:

Configuring IVR on One Switch Between Two VSANs

The following recipe provides the steps to configure IVR on a single switch. The example uses both VSANs 2112 and 3999. Before configuring IVR, you need to enable it. (See Figure 6-1.)

Figure 6-1 Single Switch IVR Topology

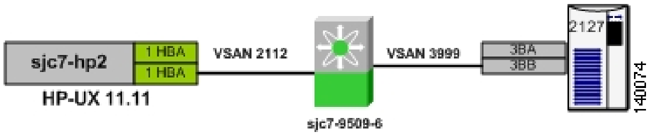

To configure IVR on one switch between two VSANs, follow these steps:

Step 1 In Fabric Manager, choose **All Vsans** > **IVR** as shown in Figure 6-2.

Figure 6-2 Enabling IVR on the Switch

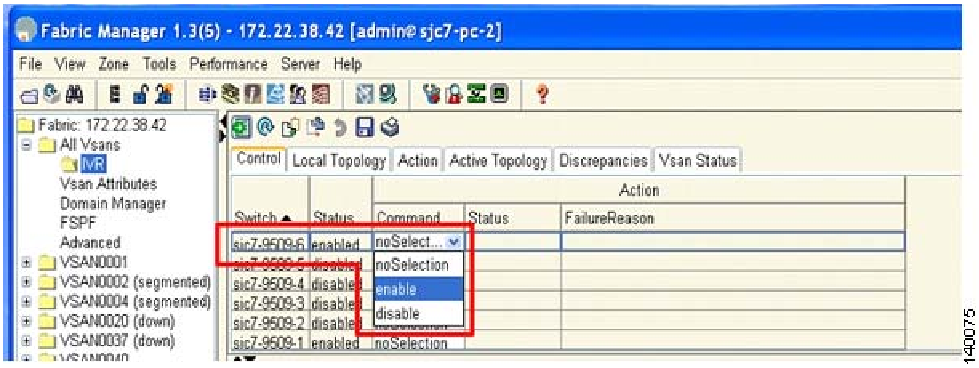

		Action		
Switch ▲	Status	Command	Status	FailureReason
sjc7-9509-6	enabled	noSelect...		
sjc7-9509-5	disabled			
sjc7-9509-4	disable			
sjc7-9509-3	disable			
sjc7-9509-2	disable			
sjc7-9509-1	enabled	noSelection		

Step 2 Choose the switch where you want to enable IVR and then enable it.

Step 3 Click the **Apply Changes** icon to save the changes as shown in Figure 6-3.

Figure 6-3 IVR Enable Status on the Switch

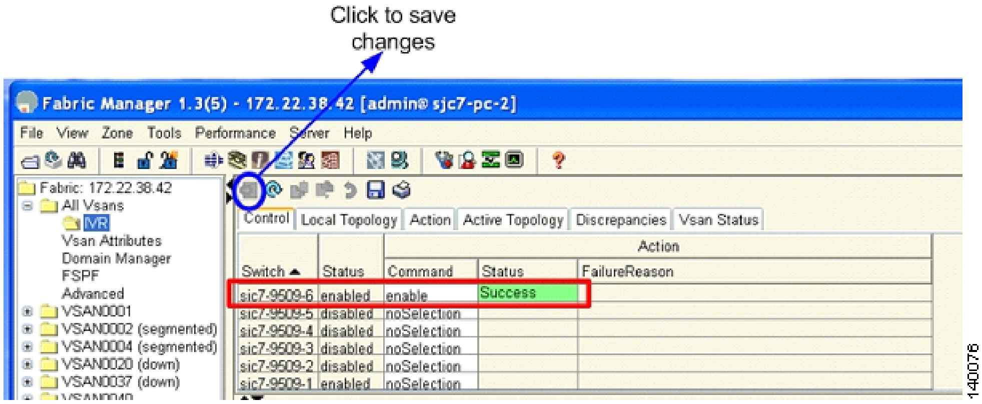
Click to save changes

		Action		
Switch ▲	Status	Command	Status	FailureReason
sjc7-9509-6	enabled	enable	Success	
sjc7-9509-5	disabled	noSelection		
sjc7-9509-4	disabled	noSelection		
sjc7-9509-3	disabled	noSelection		
sjc7-9509-2	disabled	noSelection		
sjc7-9509-1	enabled	noSelection		

Alternatively, you can enable IVR from the switch by using the **enable IVR** command.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Sjc7-9509-6(config)# ivr enable
sjc7-9509-6(config)#^Z
sjc7-9509-6#
```

Step 4 Choose the **Local Topology** tab in Fabric Manager to configure the local IVR topology.

Step 5 Click the **Create Row** icon as shown in Figure 6-4. You see a window that lists all the switches that have IVR enabled in the topology.

Figure 6-4 Inserting IVR Topology

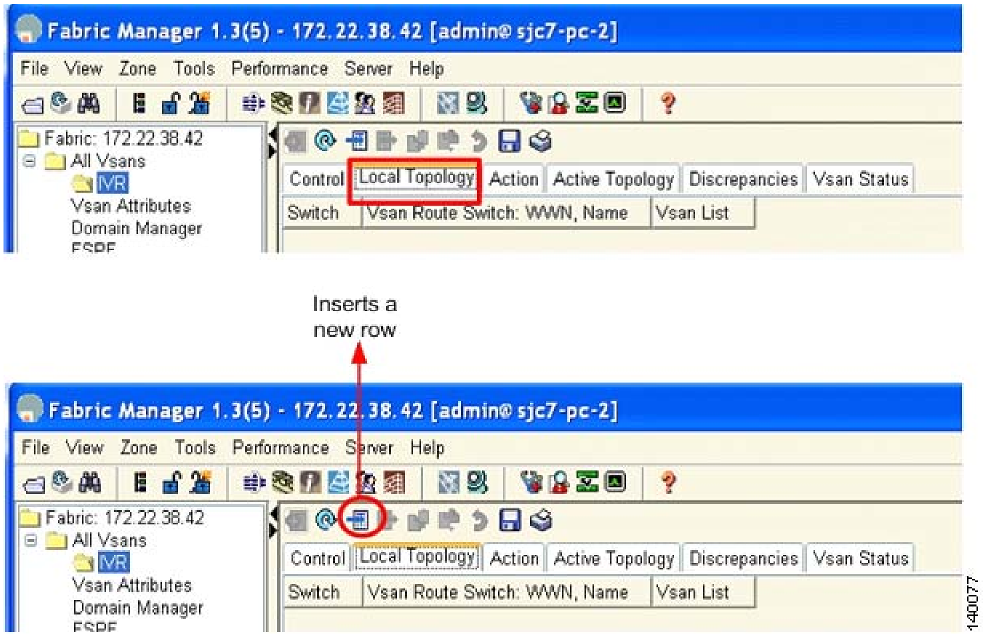

Step 6 Click the **Switch WWN** tab.

Step 7 Choose the switch WWN where you want to configure IVR.

Step 8 Choose the VSANs to add to this IVR configuration.

Figure 6-5 shows the switch sjc7-9509-6 selected, its switch WWN, and the VSANs (2112 and 3999) involved in the IVR configuration.

Figure 6-5 Adding Switch WWN and VSANs to Create the IVR Topology

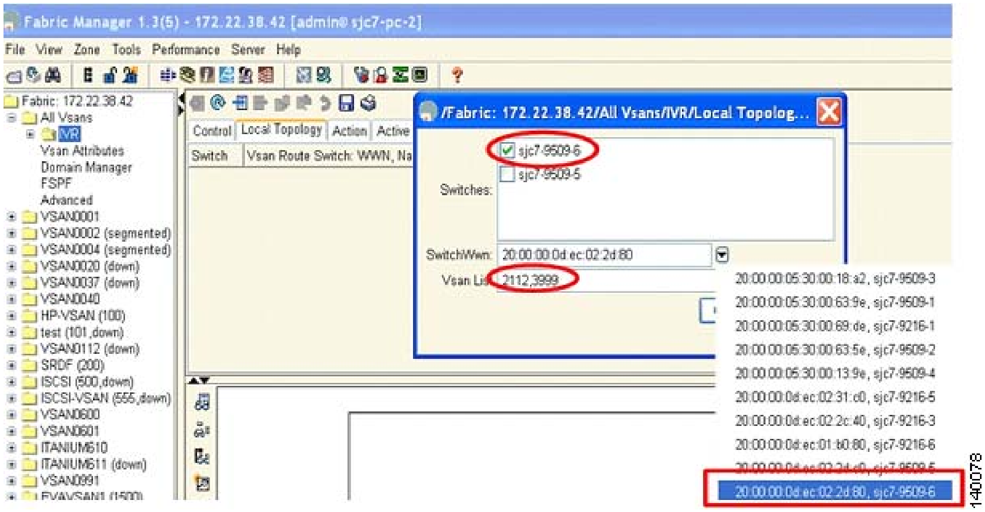

Switches: ☑ sjc7-9509-6 ☐ sjc7-9509-5

SwitchWwn: 20:00:00:0d:ec:02:2d:80

Vsan List: 2112,3999

20:00:00:05:30:00:18:a2, sjc7-9509-3
20:00:00:05:30:00:63:9e, sjc7-9509-1
20:00:00:05:30:00:69:de, sjc7-9216-1
20:00:00:05:30:00:63:5e, sjc7-9509-2
20:00:00:05:30:00:13:9e, sjc7-9509-4
20:00:00:0d:ec:02:31:c0, sjc7-9216-5
20:00:00:0d:ec:02:2c:40, sjc7-9216-3
20:00:00:0d:ec:01:b0:80, sjc7-9216-6
20:00:00:0d:ec:02:2d:c0, sjc7-9509-5
20:00:00:0d:ec:02:2d:80, sjc7-9509-6

After making these selections, the configuration is updated. You now need to activate the IVR topology.

Step 9 Activate the created IVR topology by clicking the **Action** tab.

Step 10 Check **Activate Local** to activate the created IVR topology for the switch. (See Figure 6-6.)

Figure 6-6 Activating IVR Topology

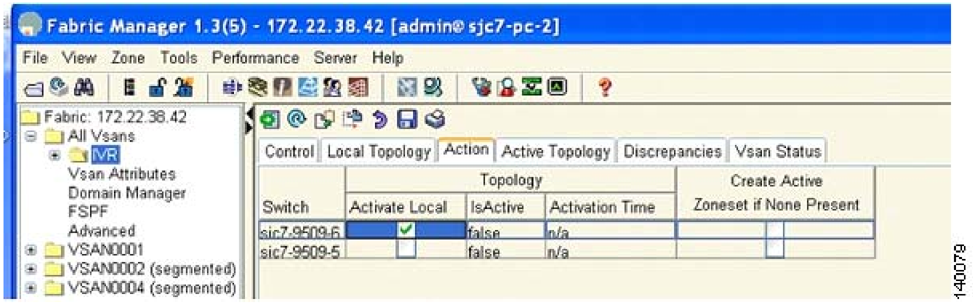

	Topology			Create Active
Switch	Activate Local	IsActive	Activation Time	Zoneset if None Present
sjc7-9509-6	☑	false	n/a	☐
sjc7-9509-5	☐	false	n/a	☐

By selecting the **Active Topology** tab you can see the active topology. The **Discrepancies** tab shows the discrepancies, if any, that need to be fixed for IVR to function properly. Ideally the discrepancies tab should not have any entries. Figure 6-7 shows the **Active Topology** tab.

Figure 6-7 Active Topology Tab

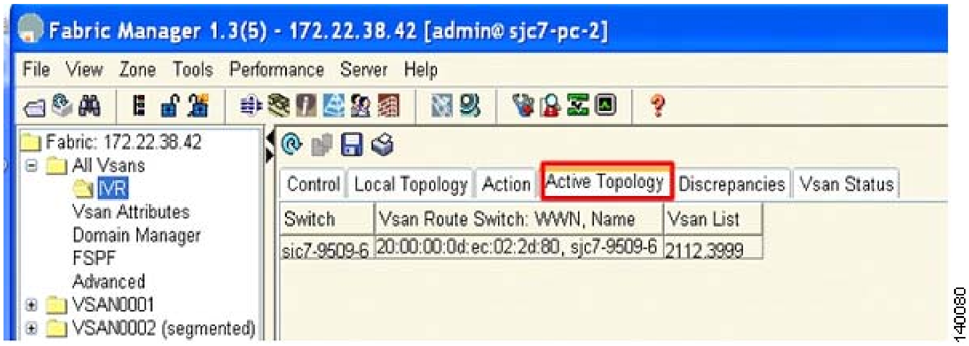

Switch	Vsan Route Switch: WWN, Name	Vsan List
sjc7-9509-6	20:00:00:0d:ec:02:2d:80, sjc7-9509-6	2112,3999

The next step is to create an IVR zone set and zones so that the initiator and the target can communicate with each other.

Step 11 Choose **IVR** > **Edit Local Full Zone Database** as shown in Figure 6-8. You see the Zone Creation window.

Figure 6-8 Edit the IVR Zone on the Switch

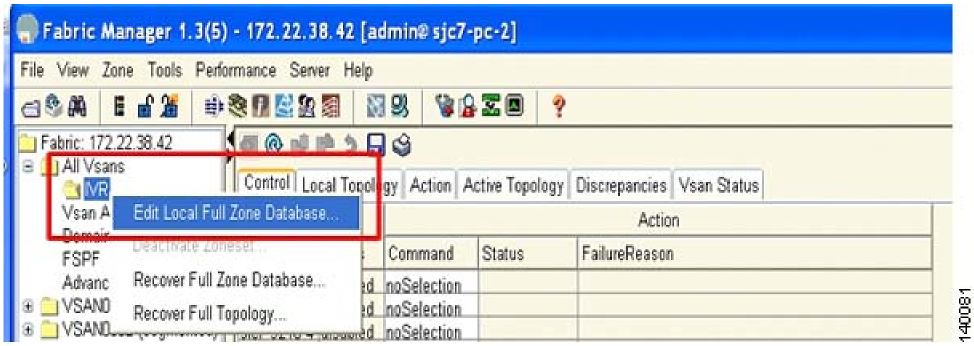

Your next step is to create a zone and assign the required members to the zone.

Step 12 Right click **Zone** and choose **Insert** as shown in Figure 6-9. You see the IVR Zone Name dialog box.

Figure 6-9 Insert IVR Zones

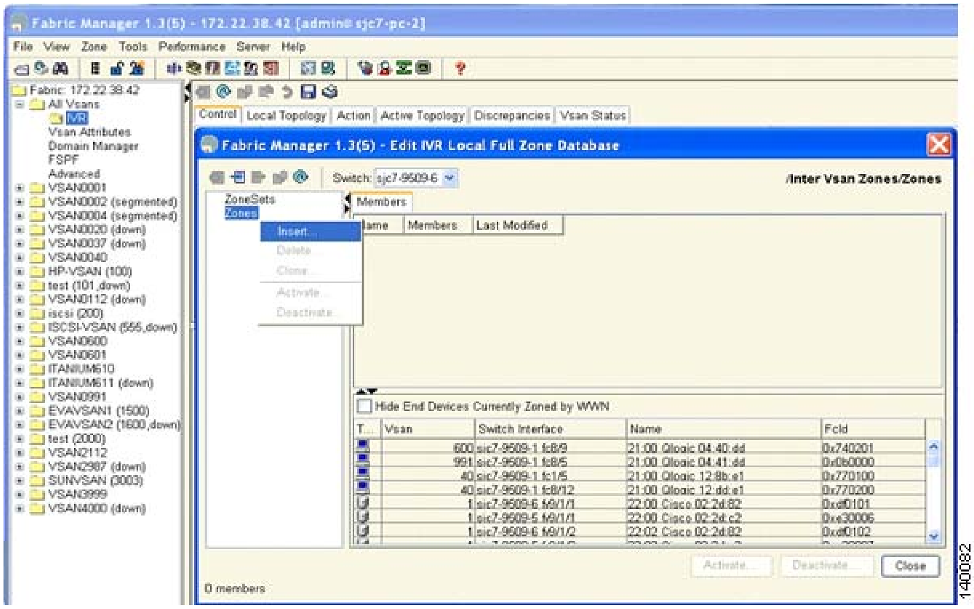

□ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	Fcld
	600	sjc7-9509-1 fc8/9	21:00 Qlogic 04:40:dd	0x740201
	991	sjc7-9509-1 fc8/5	21:00 Qlogic 04:41:dd	0x0b0000
	40	sjc7-9509-1 fc1/5	21:00 Qlogic 12:8b:e1	0x770100
	40	sjc7-9509-1 fc8/12	21:00 Qlogic 12:dd:e1	0x770200
	1	sjc7-9509-6 fv9/1/1	22:00 Cisco 02:2d:82	0xdf0101
	1	sjc7-9509-5 fv9/1/1	22:00 Cisco 02:2d:c2	0xe30006
	1	sjc7-9509-6 fv9/1/2	22:02 Cisco 02:2d:82	0xdf0102

Step 13 Enter a name for the zone, making sure you keep the prefix IVR at the beginning of the zone name. In this example, the name is IVR_Z_sjc7-hp2-0-1-1-0. (See Figure 6-10.)

Figure 6-10 Naming an IVR Zone

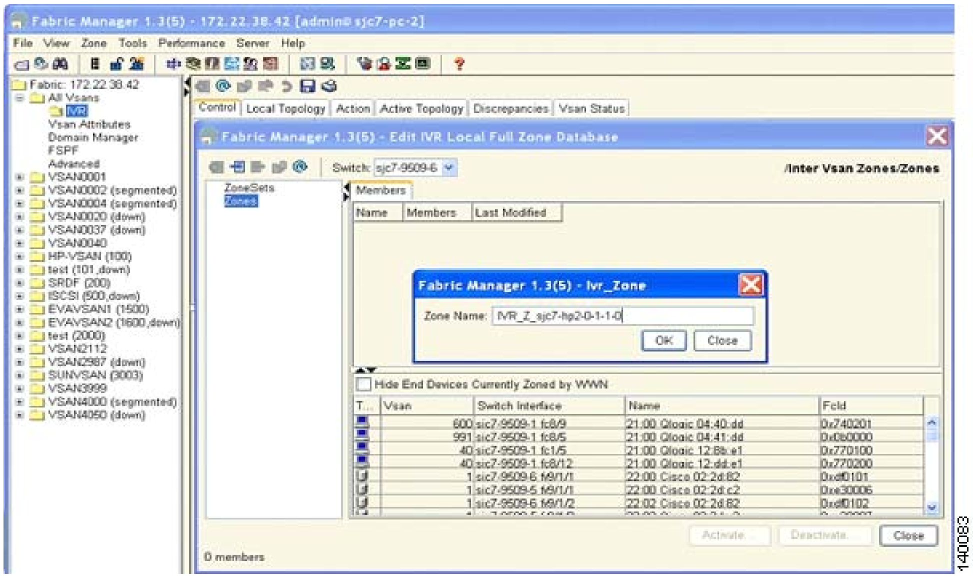

Zone Name: IVR_Z_sjc7-hp2-0-1-1-0

□ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	Fcld
	600	sjc7-9509-1 fc8/9	21:00 Qlogic 04:40:dd	0x740201
	991	sjc7-9509-1 fc8/5	21:00 Qlogic 04:41:dd	0x0b0000
	40	sjc7-9509-1 fc1/5	21:00 Qlogic 12:8b:e1	0x770100
	40	sjc7-9509-1 fc8/12	21:00 Qlogic 12:dd:e1	0x770200
	1	sjc7-9509-6 fv9/1/1	22:00 Cisco 02:2d:82	0xdf0101
	1	sjc7-9509-5 fv9/1/1	22:00 Cisco 02:2d:c2	0xe30006
	1	sjc7-9509-6 fv9/1/2	22:02 Cisco 02:2d:82	0xdf0102

Step 14 Add the member ports into zone IVR_Z_sjc7-hp2-0-1-1-0. The target port is in VSAN 3999 and the initiator port is in VSAN 2112. (See Figure 6-11.)

Figure 6-11 Members of the IVR Zone

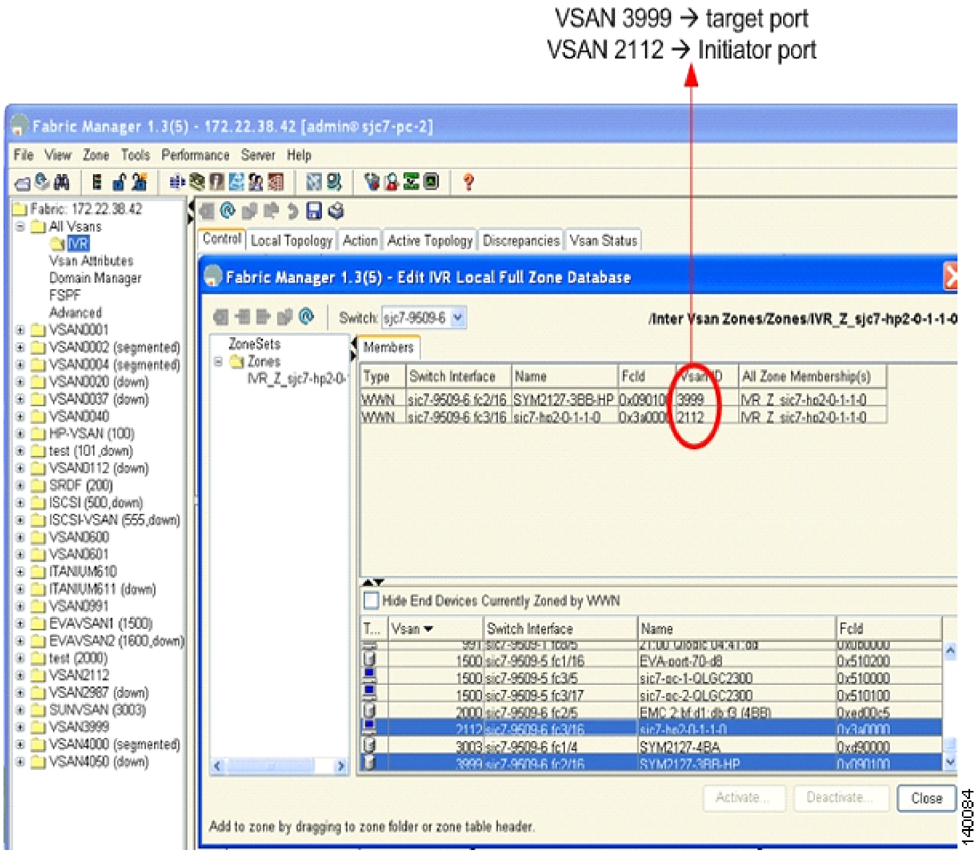

VSAN 3999 → target port
VSAN 2112 → Initiator port

Type	Switch Interface	Name	FcId	Vsan ID	All Zone Membership(s)
WWN	sjc7-9509-6 fc2/16	SYM2127-3BB-HP	0x090100	3999	IVR_Z_sjc7-hp2-0-1-1-0
WWN	sjc7-9509-6 fc3/16	sjc7-hp2-0-1-1-0	0x3a0000	2112	IVR_Z_sjc7-hp2-0-1-1-0

□ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	FcId
	991	sjc7-9509-1 fc8/5	21:00 Qlogic 04:41:dd	0x060000
	1500	sjc7-9509-5 fc1/16	EVA-port-70-d8	0x510200
	1500	sjc7-9509-5 fc3/5	sjc7-pc-1-QLGC2300	0x510000
	1500	sjc7-9509-5 fc3/17	sjc7-pc-2-QLGC2300	0x510100
	2000	sjc7-9509-6 fc2/5	EMC 2:bf:d1:db:f3 (4BB)	0xed00e5
	2112	sjc7-9509-6 fc3/16	sjc7-hp2-0-1-1-0	0x3a0000
	3003	sjc7-9509-6 fc1/4	SYM2127-4BA	0xd90000
	3999	sjc7-9509-6 fc2/16	SYM2127-3BB-HP	0x090100

Add to zone by dragging to zone folder or zone table header.

The next step is to create a IVR zone set and add the above zone to the zone set.

Step 15 In the **Edit IVR Local Full Database** window, choose **ZoneSets > Insert**. (See Figure 6-12.)

Figure 6-12 Inserting IVR Zone Sets

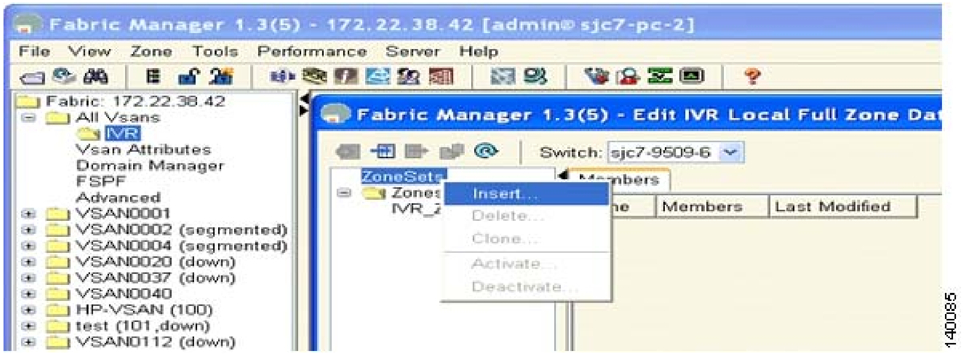

Step 16 Name the zone set, making sure to retain the IVR prefix to designate it as an IVR zone set, as shown in Figure 6-13.

Figure 6-13 Naming an IVR

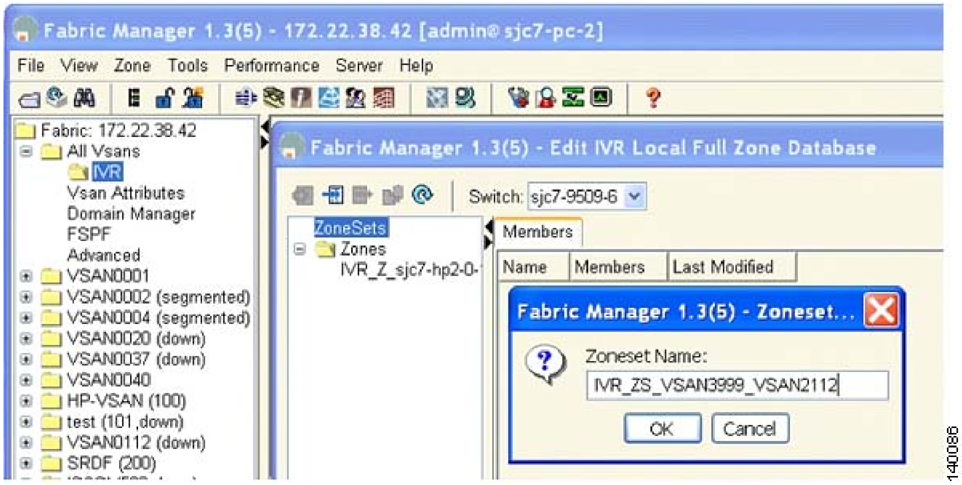

Switch: sjc7-9509-6

Zoneset Name: IVR_ZS_VSAN3999_VSAN2112

Step 17 Activate the zone set by clicking **Activate**, as shown in Figure 6-14.

Figure 6-14 Activating an IVR

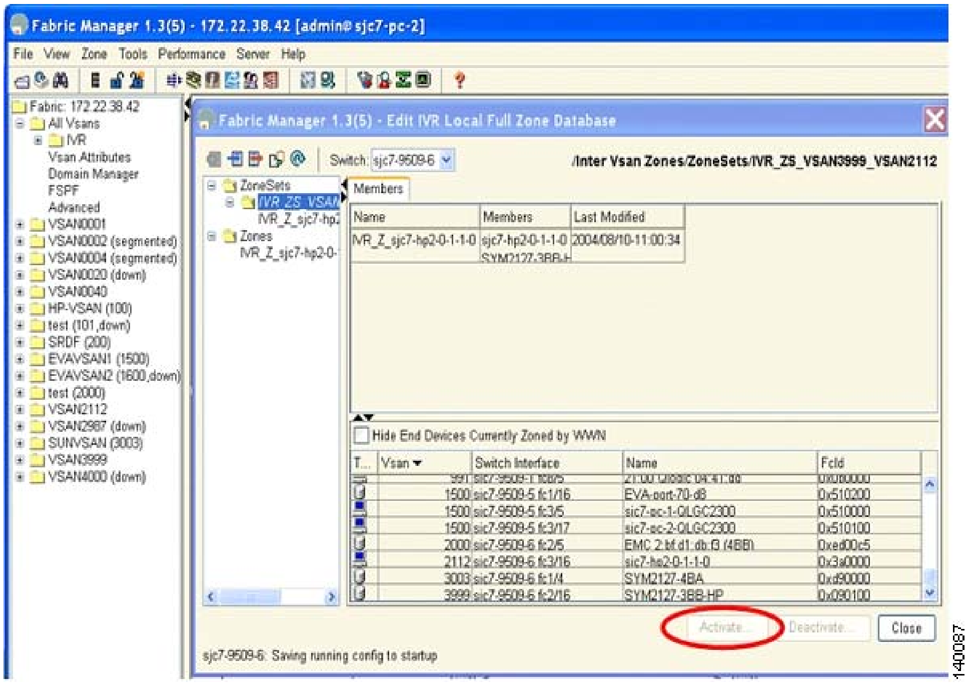

Switch: sjc7-9509-6 /Inter Vsan Zones/ZoneSets/IVR_ZS_VSAN3999_VSAN2112

Name	Members	Last Modified
IVR_Z_sjc7-hp2-0-1-1-0	sjc7-hp2-0-1-1-0 SYM2127-3BB-H	2004/08/10-11:00:34

□ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	FcId
	1500	sjc7-9509-5 fc1/16	EVA-port-70-d8	0x510200
	1500	sjc7-9509-5 fc3/5	sjc7-pc-1-QLGC2300	0x510000
	1500	sjc7-9509-5 fc3/17	sjc7-pc-2-QLGC2300	0x510100
	2000	sjc7-9509-6 fc2/5	EMC 2:bf:d1:db:f3 (4BB)	0xed00c5
	2112	sjc7-9509-6 fc3/16	sjc7-hp2-0-1-1-0	0x3a0000
	3003	sjc7-9509-6 fc1/4	SYM2127-4BA	0xd90000
	3999	sjc7-9509-6 fc2/16	SYM2127-3BB-HP	0x090100

sjc7-9509-6: Saving running config to startup

Step 18 Verify that the target and initiator can see each other by running the appropriate host commands to check that the host can see the LUNs through the target.

Configuring IVR Between Two Switches Using a Transit VSAN

When two or more switches are involved in an IVR configuration, one or more transit VSANs may be required to support IVR.

The following recipe uses switches sjc7-9509-6 and sjc7-9509-5. The target (storage) port is on switch sjc7-9509-6 in VSAN 3999 and the initiator is an HP-UX server on switch sjc7-9509-5 in VSAN 3460. VSAN 3999 is only available on switch sjc7-9509-6 and VSAN 3460 is only available on switch sjc7-9509-5. The transit VSAN used is 1234. The topology appears in Figure 6-15.

Figure 6-15 IVR Topology with Two MDS Switches

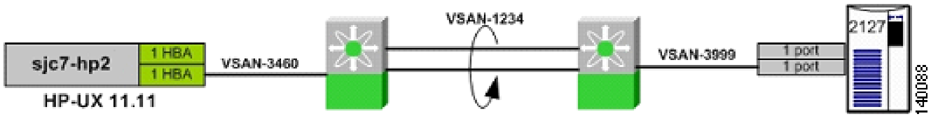

As with any IVR configuration, the switches involved must have IVR enabled on them. In this case, switches sjc7-9509-5 and sjc7-9509-6 need to have IVR enabled. Figure 6-16 shows that IVR is currently disabled on both switches

Figure 6-16 Switches with IVR Disabled

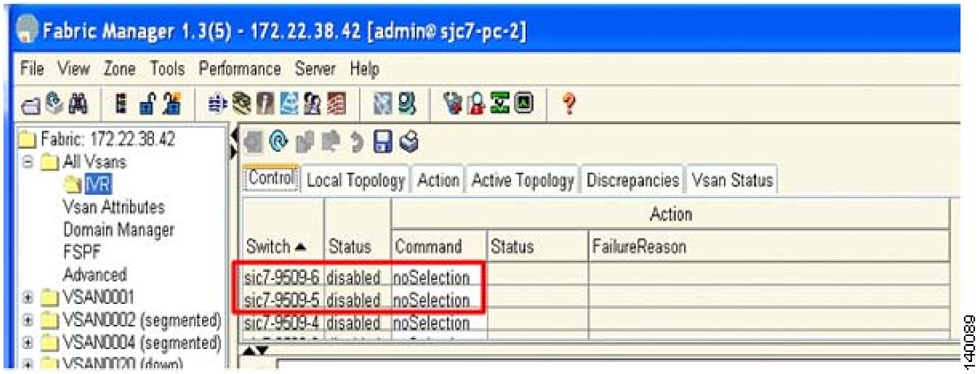

		Action		
Switch ▲	Status	Command	Status	FailureReason
sic7-9509-6	disabled	noSelection		
sic7-9509-5	disabled	noSelection		
sic7-9509-4	disabled	noSelection		

To configure IVR on two switches, follow these steps:

Step 1 In Fabric Manager, choose **All VSANs** > **IVR**, as shown in Figure 6-17.

Step 2 From the **command** pull-down menu, choose **enable** to enable IVR for a given switch.

Figure 6-17 Enabling IVR in sjc7-9509-6

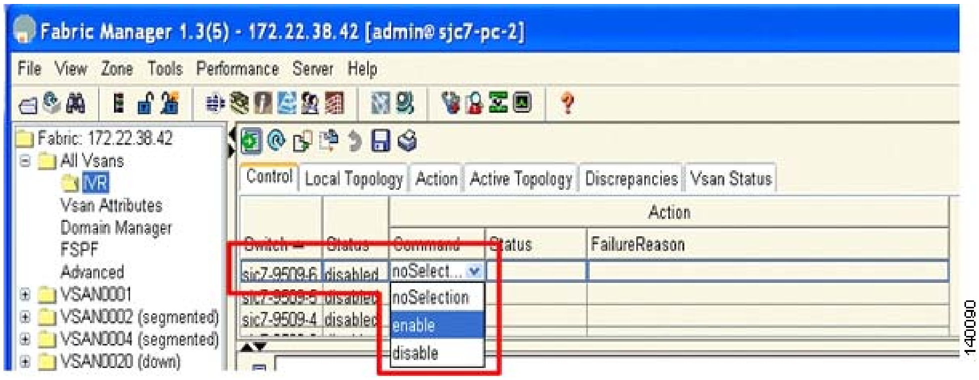

		Action		
Switch	Status	Command	Status	FailureReason
sjc7-9509-6	disabled	noSelect...		
sjc7-9509-5	disabled			
sjc7-9509-4	disabled			

Figure 6-17 shows IVR being enabled for switch sjc7-9509-6. Follow the same procedure to enable IVR for switch sjc7-9509-5.

Step 3 Click the **Apply Changes** icon to save the changes. (See Figure 6-18.)

Figure 6-18 Saving IVR Changes

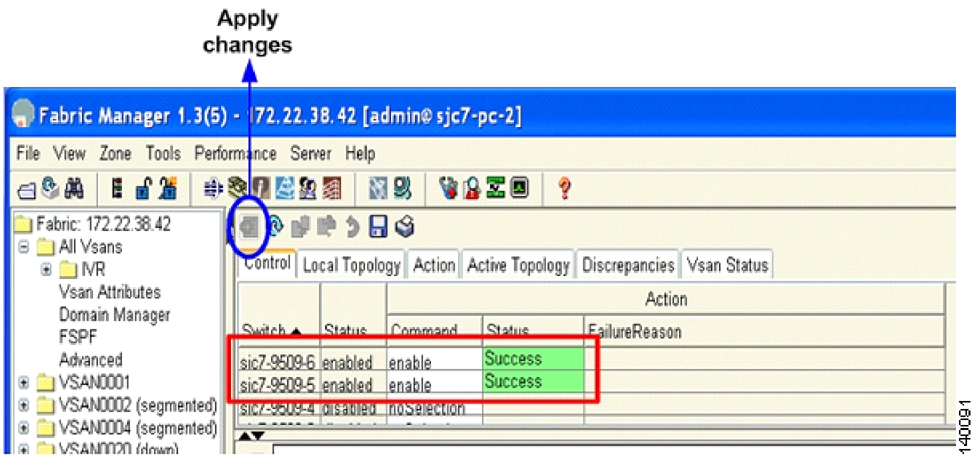

		Action		
Switch	Status	Command	Status	FailureReason
sjc7-9509-6	enabled	enable	Success	
sjc7-9509-5	enabled	enable	Success	
sjc7-9509-4	disabled	noSelection		

The next step is to create the IVR topology.

Step 4 Click the **Local Topology** tab in Fabric Manager.

Step 5 Click the **Create Row** icon, as shown in Figure 6-19. You see a window that lists all the switches that have IVR enabled in the topology

Figure 6-19 Create IVR Topology

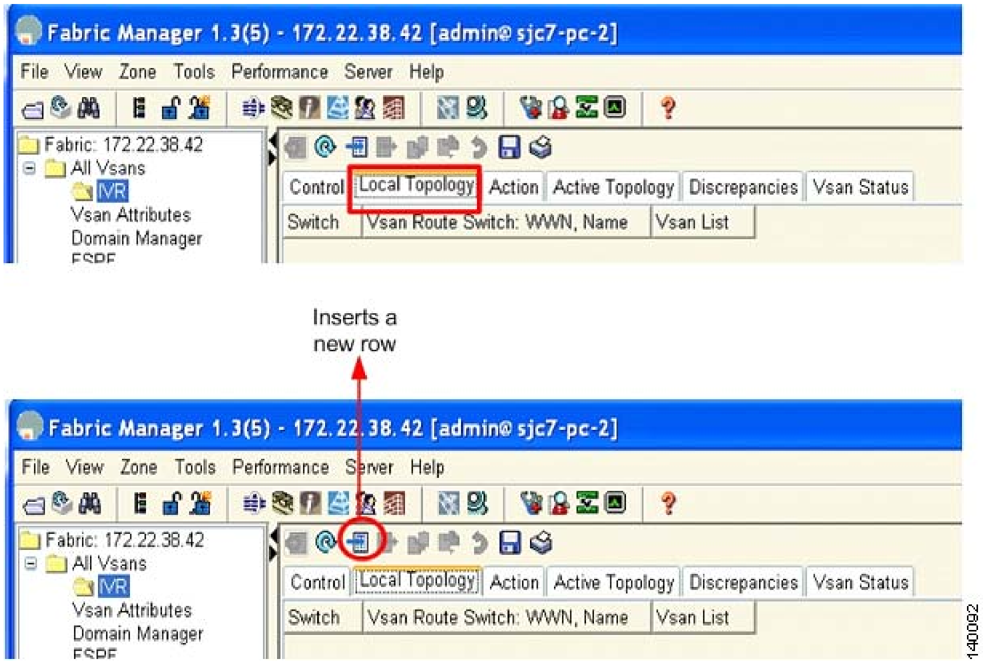

Step 6 Click the **Switch WWN** tab.

Step 7 Choose the switch WWN where you want to configure IVR.

Figure 6-20 shows that the switches sjc7-9509-5 and sjc7-9509-6 are both selected in the Switches dialog box. In the Switch WWN dialog box, the WWN of switch sjc7-9509-5 is selected. In the VSAN list, the VSANs that need to be added to complete the IVR topology are shown. In this case it is VSAN 1234 (transit VSAN) and 3460 (host VSAN). This will create the first half of the topology required to make IVR work.

Figure 6-20 Phase One of the IVR Topology Configuration

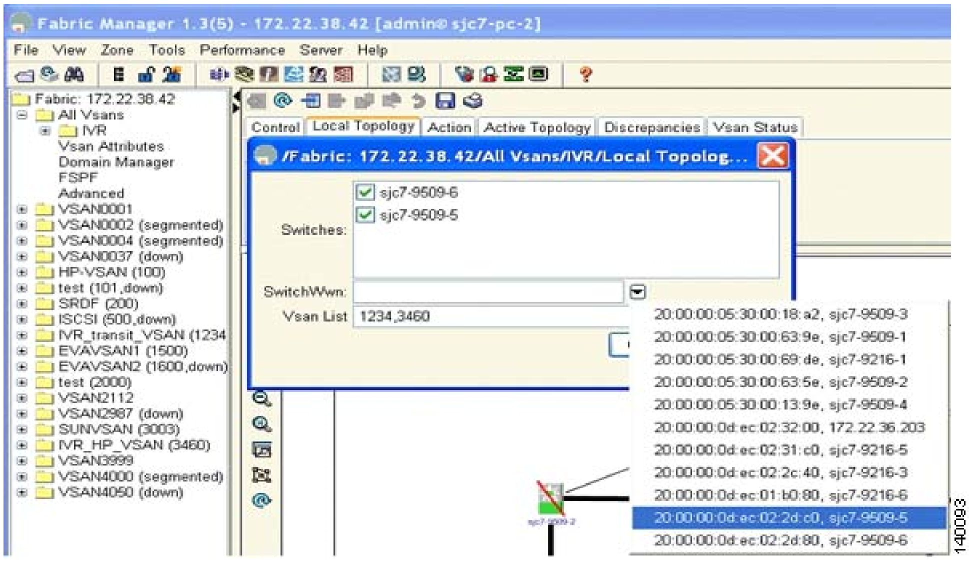

Switches: ☑ sjc7-9509-6
☑ sjc7-9509-5

SwitchWwn:

Vsan List 1234,3460

20:00:00:05:30:00:18:a2, sjc7-9509-3
20:00:00:05:30:00:63:9e, sjc7-9509-1
20:00:00:05:30:00:69:de, sjc7-9216-1
20:00:00:05:30:00:63:5e, sjc7-9509-2
20:00:00:05:30:00:13:9e, sjc7-9509-4
20:00:00:0d:ec:02:32:00, 172.22.36.203
20:00:00:0d:ec:02:31:c0, sjc7-9216-5
20:00:00:0d:ec:02:2c:40, sjc7-9216-3
20:00:00:0d:ec:01:b0:80, sjc7-9216-6
20:00:00:0d:ec:02:2d:c0, sjc7-9509-5
20:00:00:0d:ec:02:2d:80, sjc7-9509-6

Step 8 Add the VSANs on switch sjc7-9509-6 to the topology to enable IVR. This action adds both switches, the WWN of sjc7-9509-6 and VSANs 1234 (transit VSAN) and 3999 (storage VSAN) to the topology. Figure 6-21 shows the second phase of the IVR topology configuration.

Figure 6-21 Phase Two of the IVR Topology Configuration

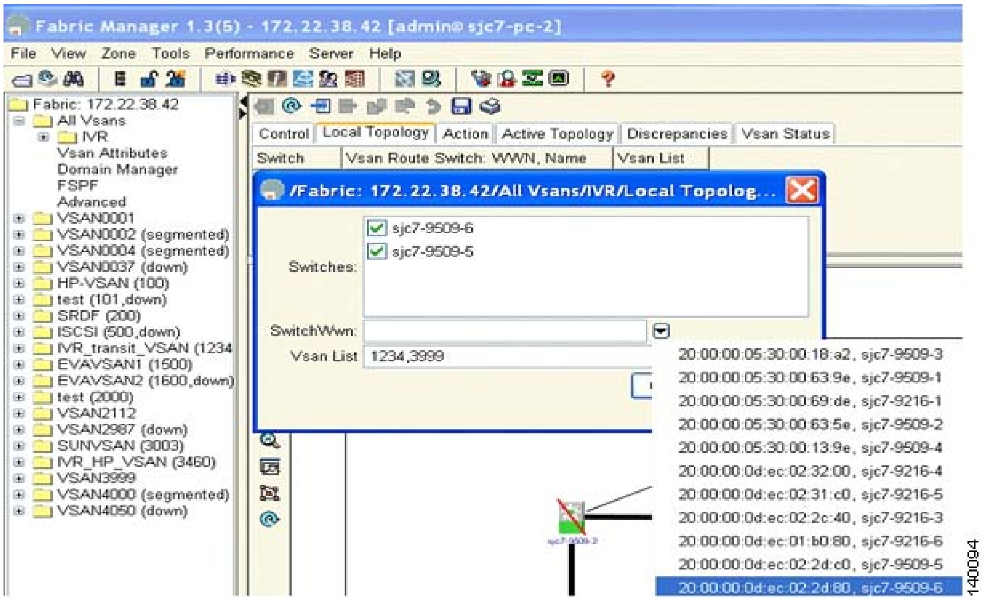

/Fabric: 172.22.38.42/All Vsans/IVR/Local Topolog...

Switches: ☑ sjc7-9509-6
☑ sjc7-9509-5

SwitchWwn:

Vsan List: 1234,3999

20:00:00:05:30:00:18:a2, sjc7-9509-3
20:00:00:05:30:00:63:9e, sjc7-9509-1
20:00:00:05:30:00:69:de, sjc7-9216-1
20:00:00:05:30:00:63:5e, sjc7-9509-2
20:00:00:05:30:00:13:9e, sjc7-9509-4
20:00:00:0d:ec:02:32:00, sjc7-9216-4
20:00:00:0d:ec:02:31:c0, sjc7-9216-5
20:00:00:0d:ec:02:2c:40, sjc7-9216-3
20:00:00:0d:ec:01:b0:80, sjc7-9216-6
20:00:00:0d:ec:02:2d:c0, sjc7-9509-5
20:00:00:0d:ec:02:2d:80, sjc7-9509-6

After the configuration, the local topology is shown in Figure 6-22. As can be seen, both switches have knowledge of all the VSANs involved in the IVR configuration.

Figure 6-22 Configured IVR Local Topology

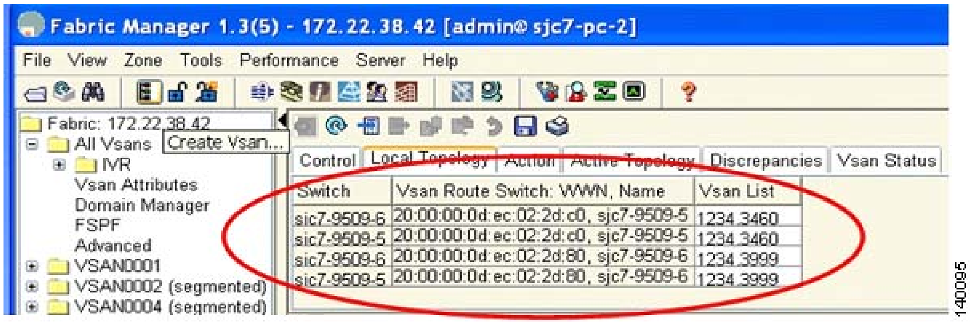

Switch	Vsan Route Switch: WWN, Name	Vsan List
sic7-9509-6	20:00:00:0d:ec:02:2d:c0, sjc7-9509-5	1234.3460
sic7-9509-5	20:00:00:0d:ec:02:2d:c0, sjc7-9509-5	1234.3460
sic7-9509-6	20:00:00:0d:ec:02:2d:80, sjc7-9509-6	1234.3999
sic7-9509-5	20:00:00:0d:ec:02:2d:80, sjc7-9509-6	1234.3999

The next step is to activate the configured IVR topology.

Step 9 In Fabric Manager, click the **Action** tab.

Step 10 Check the **Create Active Zone Set if None Present** check boxes for both switches, as shown in Figure 6-23.

Figure 6-23 Activating IVR Topology

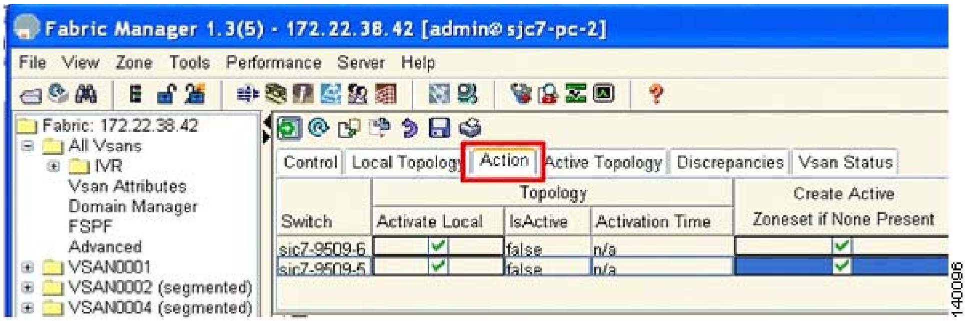

Switch	Topology Activate Local	IsActive	Activation Time	Create Active Zoneset if None Present
sjc7-9509-6	✓	false	n/a	✓
sjc7-9509-5	✓	false	n/a	✓

Tip If any of the VSANs does not have an active zone set, then the IVR zone set activation will fail if the **Create Active Zone Set if None Present** check box is not checked.

The next step is to create a IVR zone set and zones so that the initiator and the target can communicate with each other.

Step 11 Choose **IVR** > **Edit Local Full Zone Database**. You can see the zone creation window. From this window, you can create a zone and add the required members to the zone. (See Figure 6-24.)

Figure 6-24 Edit IVR Zone Set on the Switch

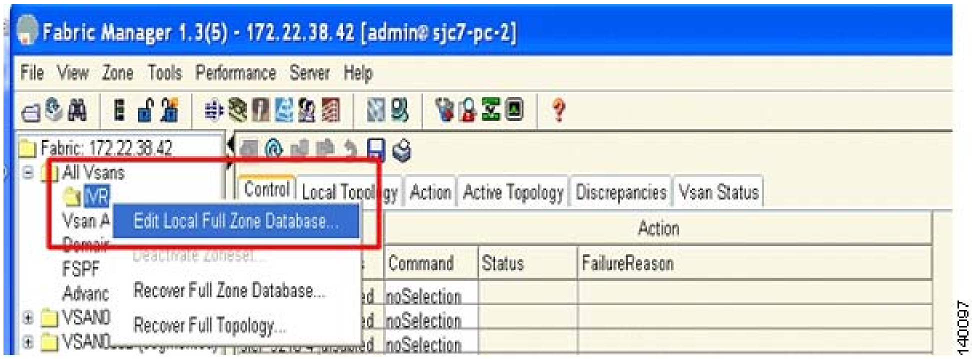

Action Command	Status	FailureReason
noSelection		
noSelection		
noSelection		

Step 12 To add a zone, right click **Zone > Insert**. You can see the IVR zone name dialog box. (See Figure 6-25.)

Figure 6-25 Insert IVR Zones

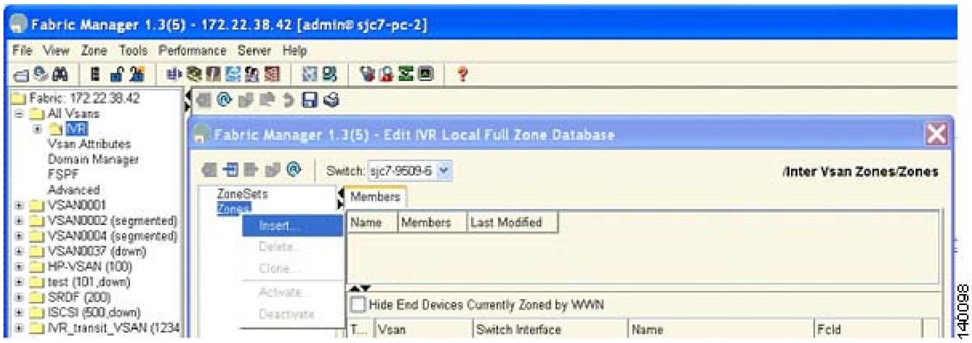

Switch: sjc7-9509-6

Name	Members	Last Modified

☐ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	Fcld

The next step is to give the zone a name.

Step 13 Enter the name IVR_Z_sjc7-hp2-0-1-3-0 for this IVR zone. The IVR prefix at the beginning of the name helps to identify the zone as an IVR zone. The name also contains the host name and the HBA hardware location of the system. (See Figure 6-26.)

Figure 6-26 Naming an IVR Zone

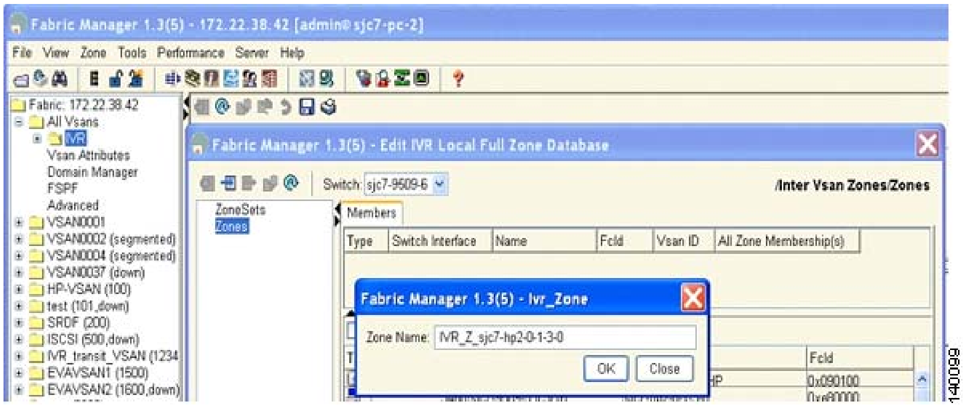

Switch: sjc7-9509-6

Zone Name: IVR_Z_sjc7-hp2-0-1-3-0

Step 14 Add the initiator and the target to the zone. This action enables them to communicate with each other. (See Figure 6-27.)

Figure 6-27 Members of the IVR Zone

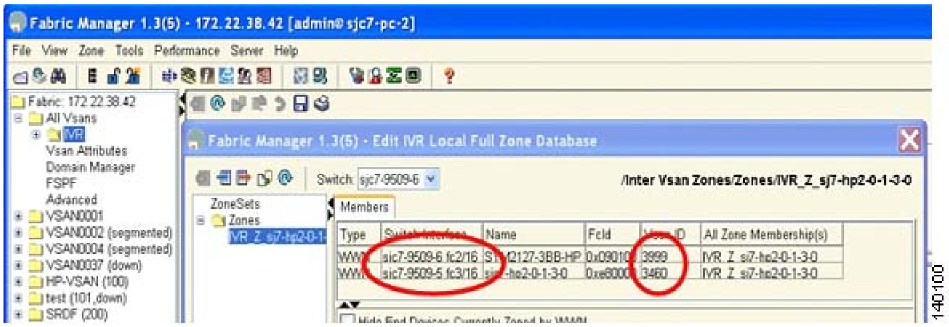

Switch: sjc7-9509-6

/Inter Vsan Zones/Zones/IVR_Z_sj7-hp2-0-1-3-0

Type	Switch Interface	Name	Fcld	Vsan ID	All Zone Membership(s)
WW	sjc7-9509-6 fc2/16	S[illegible]2127-3BB-HP	0x09010	3999	IVR_Z_sj7-hp2-0-1-3-0
WW	sjc7-9509-5 fc3/16	sj[illegible]-hp2-0-1-3-0	0xe8000	3460	IVR_Z_sj7-hp2-0-1-3-0

The next step is to create the IVR zone set and activate the zone set to complete the IVR configuration.

Step 15 In the Edit IVR Local Full Database window, choose **zone sets > Insert**. (See Figure 6-28.)

Figure 6-28 Inserting IVR Zone Sets

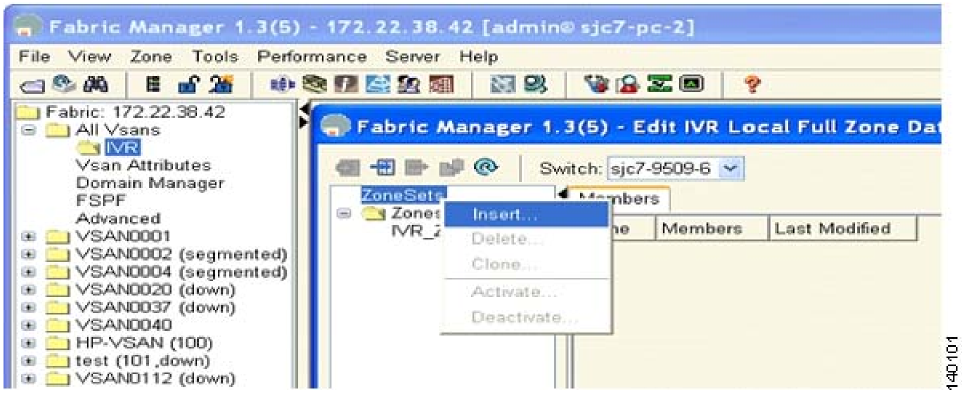

Switch: sjc7-9509-6

Step 16 Add the name IVR_ZS_1234_3460_3999 for the zone set. The name IVR classifies it as an IVR zone set and contains the VSAN number to show what VSANs are involved in the configuration. (See Figure 6-29.)

Figure 6-29 IVR Zone Set

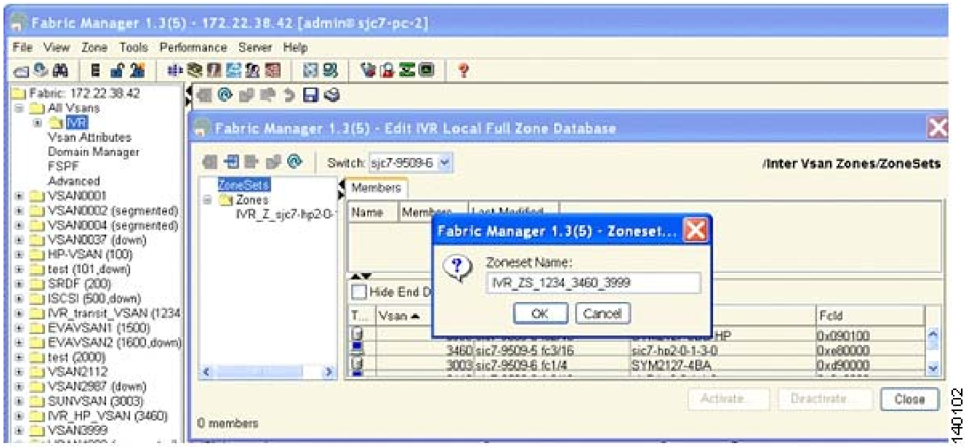

Switch: sjc7-9509-6

Zoneset Name: IVR_ZS_1234_3460_3999

T...	Vsan			Fcld
			HP	0x090100
	3460	sjc7-9509-5 fc3/16	sjc7-hp2-0-1-3-0	0xe80000
	3003	sjc7-9509-6 fc1/4	SYM2127-4BA	0xd90000

0 members

Step 17 Add the zone IVR_Z_sjc7-hp2-0-1-3-0 to the zone set. (See Figure 6-30.)

Figure 6-30 Adding an IVR Zone to an IVR Zone Set

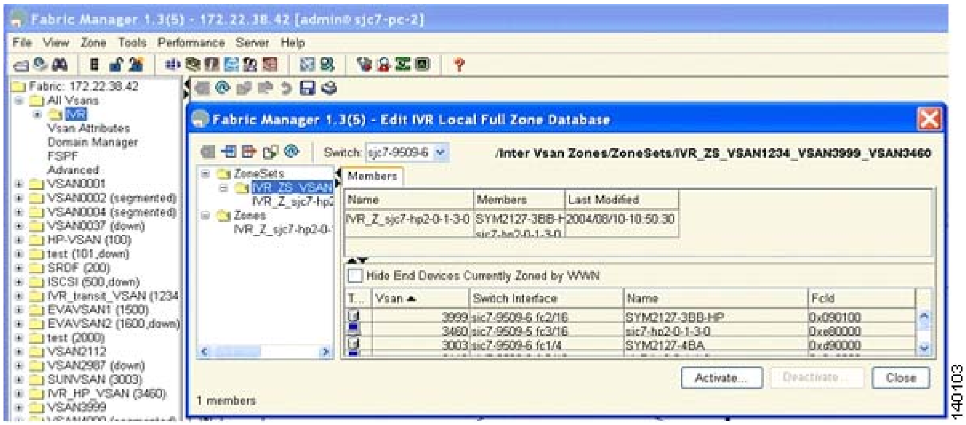

Fabric Manager 1.3(5) - 172.22.38.42 [admin@sjc7-pc-2]

Fabric Manager 1.3(5) - Edit IVR Local Full Zone Database

Switch: sjc7-9509-6

/Inter Vsan Zones/ZoneSets/IVR_ZS_VSAN1234_VSAN3999_VSAN3460

Name	Members	Last Modified
IVR_Z_sjc7-hp2-0-1-3-0	SYM2127-3BB-H sjc7-hp2-0-1-3-0	2004/08/10-10:50:30

☐ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	FcId
	3999	sjc7-9509-6 fc2/16	SYM2127-3BB-HP	0x090100
	3460	sjc7-9509-5 fc3/16	sjc7-hp2-0-1-3-0	0xe80000
	3003	sjc7-9509-6 fc1/4	SYM2127-4BA	0xd90000

1 members

Step 18 Activate the zone set by clicking **Activate**, as shown in Figure 6-31.

Figure 6-31 Activating an IVR Zone Set

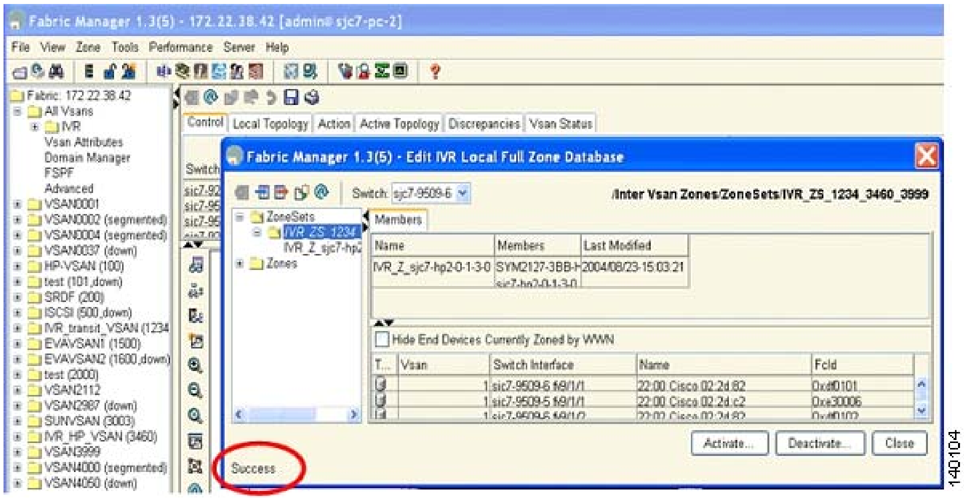

Fabric Manager 1.3(5) - 172.22.38.42 [admin@sjc7-pc-2]

Fabric Manager 1.3(5) - Edit IVR Local Full Zone Database

Switch: sjc7-9509-6

/Inter Vsan Zones/ZoneSets/IVR_ZS_1234_3460_3999

Name	Members	Last Modified
IVR_Z_sjc7-hp2-0-1-3-0	SYM2127-3BB-H sjc7-hp2-0-1-3-0	2004/08/23-15:03:21

☐ Hide End Devices Currently Zoned by WWN

T...	Vsan	Switch Interface	Name	FcId
	1	sjc7-9509-6 fv9/1/1	22:00 Cisco 02:2d:82	0xdf0101
	1	sjc7-9509-5 fv9/1/1	22:00 Cisco 02:2d:c2	0xe30006
	1	sjc7-9509-6 fv9/1/2	22:00 Cisco 02:2d:82	0xdf0102

Success

Step 19 Verify that the target and initiator can see each other by running the appropriate host commands to check that the host can see the LUNs from the target.

The two-switch IVR configuration is now complete.

IVR with Brocade and McData Switches Using Interop Mode

The procedure to configure IVR in interop mode with third-party switches is the same as described in the preceding recipes. You can use IVR in conjunction with Brocade and McData switches. When a third- party switch is present in the topology, they can be located either in the transit VSAN or as an edge switch connecting to an MDS 9000 switch enabled for IVR. Third-party switches cannot be used as IVR gateways. Figure 6-32 shows topologies such as these.

Figure 6-32 Sample IVR Topology with Interop Modes

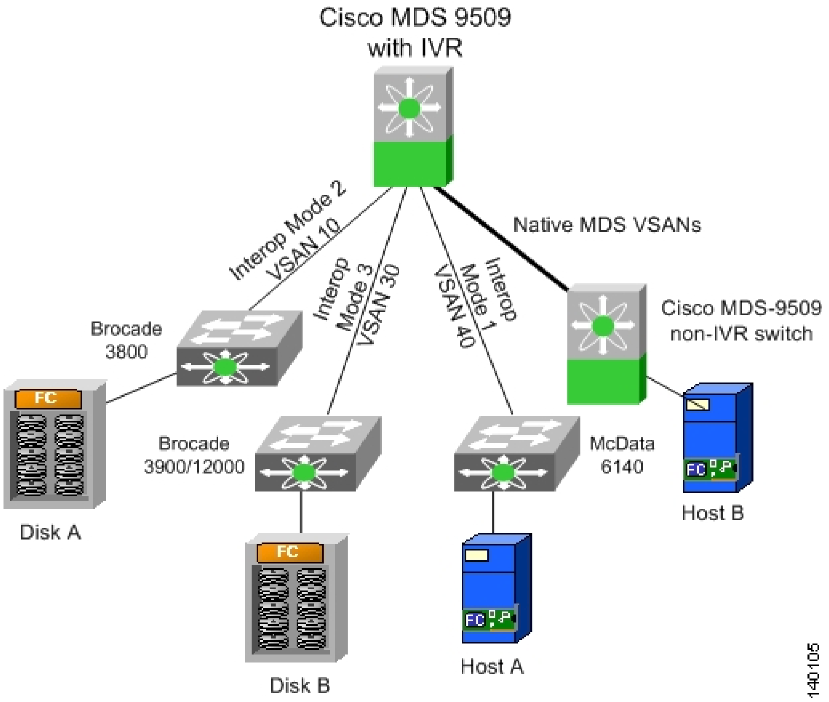

For the third-party switches to function properly in an IVR configuration, they must be configured in a VSAN interop mode. IVR works with all three interop modes (1, 2, and 3). For more details on interoperability, refer to the *MDS Switch to Switch Interoperability Configuration Guide* located on http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_mds_integration_list.html.

When a McData switch is present, the domain ID of the destination switch needs to be in the interop range. That is, the domain ID of the switch in a VSAN has to be between 97 and 127 for IVR to function properly. This means that a device on a McData switch cannot send a frame to a device in another VSAN, regardless of the interop mode, that is outside of the 97 to127 domain ID range.

CHAPTER 7

Zoning

Zones and zone sets are the basic form of data path security within a Fibre Channel environment. A zone set is a collection of zones which in turn have individual members in them. Only those members within the same zone can communicate with each other. A device can be a member of multiple zones and those devices not in a zone are in the default zone. The policy for the default zone can either be to permit devices to see each other or to deny devices in the default zone from seeing each other.

This chapter focuses on the creation of zones, zone sets, and ways to manipulate them. It includes the following sections:

Zones

In order for two devices to communicate, they must be in the same zone. Valid members of a zone can be:

- Port WWN (pWWN)
- Fibre Channel alias
- FC ID
- FWWN (WWN of a Fibre Channel interface)
- Switch interface (Fibre Channel X/Y)
- Symbolic node name

The three most common types of zone members are pWWN, FC alias, and the switch interface.

Tip We recommend that pWWN (or a Fibre Channel Alias representing a pWWN) be used for zoning as it provides the most security and ties a zone member to a specific HBA rather than to the switch port.

The name that you choose for the zone is very important. Many environments use different zone names, however, all name formats should provide relevant information as to their contents. Names like Zone1 or TapeZone do not provide sufficient information about their contents.

A zone name should contain two members and, within the zone name, contain identifiers related to the two devices, such as Z_testhost_fcaw0_symm13FA3aa. The name may be longer than Z_testhost_hba0, but should provide enough detailed information about the contents that consulting further documentation is not necessary.

Creating a Zone and Adding It to a Zone Set (Standalone Method)

This procedure demonstrates how to create a single zone with a Solaris host and a disk storage port in it, and then add it to the zone set ZS_Engr_primary. The procedure uses the standalone method, which does not automatically add the zone to the zone set upon creation of the zone. You can also use this procedure to determine how to add an existing zone to a zone set.

This example uses pWWNs as the zone members, which can be obtained either from the device itself or from the **show flogi database vsan 804** command. (See Figure 7-1.)

Resources:

Zone set: ZS_Engr_primary

Solaris host1, hba instance fcaw0: 22:35:00:0c:85:e9:d2:c2

Symmetrix 78, FA port 03ab: 10:00:00:00:c9:32:8b:a8

Figure 7-1 Standalone Zoning Topology

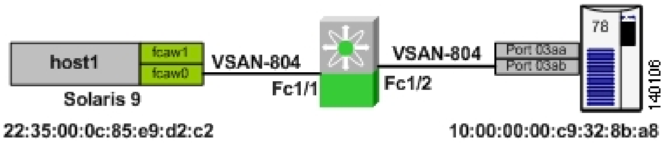

To create a single zone and add it to a zone set using the standalone method, follow these steps:

Step 1 Create the zone, building a zone name that reflects the names of the members.

```
ca-9506# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zone name Z_host1_fcaw0_symm78FA03ab vsan 804
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
```

Step 2 Add the zone to the zone set.

```
ca-9506(config)# zoneset name ZS_Engr_primary vsan 804
ca-9506(config-zoneset)# member Z_host1_fcaw0_symm78FA03ab
```

Step 3 Display the zone set.

```
ca-9506# show zoneset name ZS_Engr_primary vsan 804
zoneset name ZS_Engr_primary vsan 804
  zone name Z_host1_fcaw0_symm78FA03ab vsan 804
    pwwn 22:35:00:0c:85:e9:d2:c2
    pwwn 10:00:00:00:c9:32:8b:a8
```

Step 4 Finally, to put the zone set into production, activate the it using **zone set activate name ZS_Engr_primary vsan 804.** This command activates all the zones in the zone set, not just the one just added.

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

Creating a Zone and Adding It to a Zone Set (Inline Method)

This procedure demonstrates how to create a single zone with a Solaris host and a disk storage port in it, and add it to the zone set ZS_Engr_primary. The procedure uses the inline method, which automatically adds the zone to the zone set upon creation of the zone.

This example uses pWWNs as the zone members, which can be obtained either from the device itself or from the **show flogi database vsan 804** command. (See Figure 7-2.)

Resources:

Zone set: ZS_Engr_primary

Solaris host1, hba instance fcaw0: 22:35:00:0c:85:e9:d2:c2

Symmetrix 78, FA port 03ab: 10:00:00:00:c9:32:8b:a8

Figure 7-2 Inline Zoning Topology

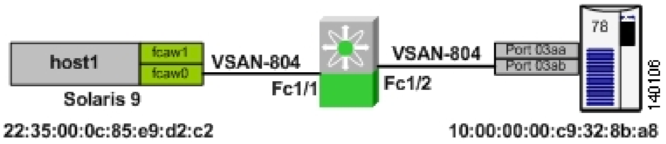

To create a single zone and add it to a zone set using the inline method, follow these steps:

Step 1 Enter the submode of the zone set.

```
ca-9506# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset name ZS_Engr_primary vsan 804
```

Step 2 Create the zone.

```
ca-9506(config-zoneset)# zone name Z_host1_fcaw0_symm78FA03ab
```

Step 3 Add the members.

```
ca-9506(config-zoneset-zone)# member pwwn 22:35:00:0c:85:e9:d2:c2
ca-9506(config-zoneset-zone)# member pwwn 10:00:00:00:c9:32:8b:a8
```

Step 4 Display the zone set.

```
ca-9506# show zoneset name ZS_Engr_primary vsan 804
zoneset name ZS_Engr_primary vsan 804
  zone name Z_host1_fcaw0_symm78FA03ab vsan 804
    pwwn 22:35:00:0c:85:e9:d2:c2
    pwwn 10:00:00:00:c9:32:8b:a8
```

Step 5 Finally, to put the zone set into production, activate the zone set using **zone set activate name ZS_Engr_primary vsan 804.** This command activates all the zones in the zone set, not just the one just added.

Creating a Fibre Channel Alias Based Zone

Fibre Channel aliases allow the administrator to assign a plain text, human readable name to a pWWN, FC ID, interface, IP address, nWWN or symbolic node name. Fibre Channel aliases are restricted to the VSAN in which they were created. The most common and recommended Fibre Channel alias is the pWWN, which is the basis for this procedure. (See Figure 7-3.)

Tip

- Aliases are distributed with the full zone set database. Therefore, if zoning is going to be edited on multiple switches, full zone set distribution should be enabled.
- An alias can be mapped to more than one device, however, we recommend that a one-to-one mapping be used.

Resources:

Zone set: ZS_Engr_primary

Solaris host1, hba instance fcaw0: 22:35:00:0c:85:e9:d2:c2

Symmetrix 78, FA port 03ab: 10:00:00:00:c9:32:8b:a8

Figure 7-3 Alias Zoning Topology

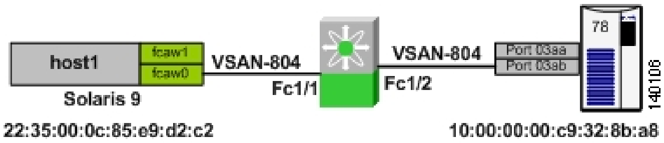

To create a Fibre Channel alias based zone, follow these steps:

Step 1 Create the Fibre Channel alias to pWWN mappings.

```
ca-9506# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# fcalias name host1_fcaw0 vsan 804
ca-9506(config-fcalias)# member pwwn 22:35:00:0c:85:e9:d2:c2
ca-9506(config-fcalias)# exit
ca-9506(config)# fcalias name symm78_fa03ab vsan 804
ca-9506(config-fcalias)# member pwwn 10:00:00:00:c9:32:8b:a8
ca-9506(config-fcalias)# end
```

Step 2 Display the newly created Fibre Channel aliases.

```
ca-9506# show fcalias vsan 804
fcalias name host1_fcaw0 vsan 804
  pwwn 22:35:00:0c:85:e9:d2:c2

fcalias name symm78_fa03ab vsan 804
  pwwn 10:00:00:00:c9:32:8b:a8
```

Step 3 Create the zone using the Fibre Channel aliases.

```
ca-9506# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset name ZS_Engr_primary vsan 804
ca-9506(config-zoneset)# zone name Z_host1_fcaw0_symm78FA03ab
ca-9506(config-zoneset-zone)# member fcalias host1_fcaw0
ca-9506(config-zoneset-zone)# member fcalias symm78_fa03ab
```

Step 4 Optionally, display the zone set.

```
ca-9506# show zoneset vsan 804
zoneset name ZS_Engr_primary vsan 804
  zone name Z_host1_fcaw0_symm78FA03ab vsan 804
    fcalias name host1_fcaw0 vsan 804
      pwwn 22:35:00:0c:85:e9:d2:c2

    fcalias name symm78_fa03ab vsan 804
      pwwn 10:00:00:00:c9:32:8b:a8
```

Step 5 Activate the zone set.

```
ca-9506# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset activate name ZS_Engr_primary vsan 804
Zoneset activation initiated. check zone status
```

Creating an Interface Based Zone

This procedure describes how to create a zone based upon the physical interface (fc X/Y) of the switch. (See Figure 7-4.)

Tip Use interface based zoning when you need to create a zone prior to connecting the HBA to the fabric. After connecting the HBA to the fabric, convert the zone member to a pWWN based member.

Figure 7-4 Interface Zoning Topology

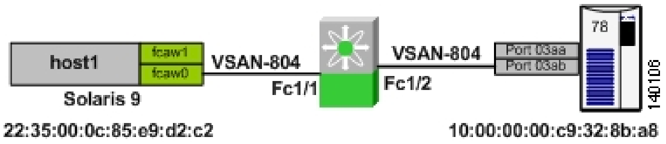

To create a zone based on the physical interface of the switch, follow these steps:

Step 1 Create the zone using the interfaces.

```
ca-9506# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset name ZS_Engr_primary vsan 804
ca-9506(config-zoneset)# zone name Z_host1_fcaw0_symm78FA03ab
ca-9506(config-zoneset-zone)# member interface fc1/1
ca-9506(config-zoneset-zone)# member interface fc1/2
```

Step 2 Optionally, display the zone set.

```
ca-9506# show zoneset vsan 804
zoneset name ZS_Engr_primary vsan 804
  zone name Z_host1_fcaw0_symm78FA03ab vsan 804
    interface fc1/1 swwn 20:00:00:0c:85:e9:d2:c0
    interface fc1/2 swwn 20:00:00:0c:85:e9:d2:c0
```

Note The sWWN is the switch's WWN, as displayed by the **show wwn switch** command:

```
ca-9506# show wwn switch
Switch WWN is 20:00:00:0c:85:e9:d2:c0
```

Step 3 Activate the zone set.

```
ca-9506# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset activate name ZS_Engr_primary vsan 804
Zoneset activation initiated. check zone status
```

Zone Sets

Zone sets are containers of zones. There are two zone set types on the Cisco MDS 9000 platform:

- **Active Zone set**—The active zone set provides the rules by which the Cisco MDS 9000 platform enforces its zoning security policy. The active zone set cannot be modified and is distributed to all switches in the VSAN. There are specific rules to merging the active zone set when two switches are connected by an ISL, as set by the Fibre Channel standards.
- **Local Zone set**—The local zone sets are contained in the full zone set database on the switch. The zone sets can be edited directly and then activated to become the active zone set. They can optionally be distributed to other switches, either manually or when a zone set is activated.

Zone Set Distribution

There are two distribution methods for zone sets: automatic and manual.

Automatic Zone Set Distribution

To enable the switch to distribute the local zone set to all other switches in the VSAN when a zone set is activated, use the **zone set distribute full vsan 804** command:

```
ca-9506# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
ca-9506(config)# zoneset distribute full vsan 804
```

Tip You can enable the automatic distribution feature on all switches in the fabric by specifying it in the initial setup script.

Manual Zone Set Distribution

To distribute the full zone set database to other switches without activating a zone set, use the **zone set distribute vsan 804** command. This method can be effective when a new switch is brought into the fabric and the zone set with its zones and Fibre Channel aliases needs to be distributed. This command will overwrite the exiting zone set database in the target switch.

```
ca-9506# zoneset distribute vsan 804
Zoneset distribution initiated. check zone status
```

CHAPTER 8

iSCSI

iSCSI is a transport protocol that is used to transport SCSI packets over TCP/IP. It is an Internet Protocol (IP) based standard that allows hosts to connect to and access storage over a network interface card. iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. Since iSCSI runs over IP, it can be used to transmit data over LAN, WAN, MAN, and so on, thereby enabling data access that is independent of the location of the storage subsystem. The Cisco MDS 9200 and 9500 series switches support iSCSI using the IPS-8 blade, IPS-4 blade, and the 14/2-port Multiprotocol Services (MPS-14/2) module.

This chapter includes the following sections:

Enabling iSCSI

Before you can configure iSCSI on the switch, you must enable it.

Warning **Do no attempt to perform any iSCSI configuration without first enabling it on the switch. You can enable iSCSI through Fabric Manager or using the iscsi enable command.**

```
MDS1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
MDS1(config)# iscsi enable
MDS1(config)#^Z
MDS1#
```

Configuring iSCSI in Transparent Mode

The following recipe shows how to configure iSCSI on a Cisco MDS 9000 switch. Transparent mode configures an equivalent Fibre Channel initiator for each iSCSI initiator. In this process, there is no LUN masking or reassignment on the MDS 9000 switch. For larger installations, configure iSCSI using the proxy initiator mode, as described in Configuring iSCSI in Proxy Initiator Mode, page 8-11.

This example uses the transparent mode to configure iSCSI on a Cisco MDS 9000 switch. The topology is shown in Figure 8-1.

Figure 8-1 iSCSI Topology

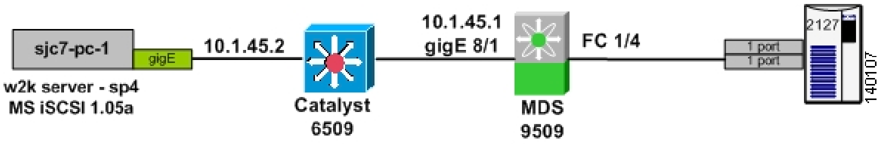

The topology consists of a Windows 2000 server with a dedicated Gigabit Ethernet NIC (network interface card) for iSCSI. It is connected to a port on the Catalyst 6509. The iSCSI interface on the host has the IP address 10.1.45.2. The IPS port 8/1 on the MDS 9509 switch is also connected to the Catalyst 6509 and has the IP address 10.1.45.1. The storage port from the array is connected to the FC port 1/4 on the MDS 9509. This example shows how to configure an iSCSI initiator using IQN (**i**SCSI **q**ualified **n**ame).

Warning **The IP address for the ports on the IPS blade should be in a different subnet and Ethernet segment than the management interface. This is critical to get iSCSI working on a MDS 9000 Family switch.**

To enable the iSCSI initiator to see drives from the array, follow these steps:

Step 1 Configure the Gigabit Ethernet interface on the MDS 9509 switch by assigning it an IP address and subnet mask. The address allows the Gigabit Ethernet interface to communicate with the network.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface gigabitethernet 8/1
sjc7-9509-6(config-if)# ip address 10.1.45.1 255.255.255.0
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 2 Configure the IP route if required.

The IP route should be configured to communicate with the initiator if the initiator is in another subnet. In the preceding step, both the initiator and the iSCSI interface on the MDS 9000 switch are in the same subnet. Therefore, the configuration of an explicit route on the MDS 9509 switch is not required.

The syntax for configuring a route is shown in the following example. Here an initiator is in the 10.1.46.0 subnet, hence a route is added to reach that subnet from the MDS 9509 switch.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# ip route 10.1.46.0 255.255.255.0 10.1.45.2
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 3 Configure the IP address on the interface on the Windows 2000 server.

Configure the IP address 10.1.45.2 on the network interface dedicated for iSCSI on the Windows 2000 server. You may need to add an IP route on the host to reach the Gigabit Ethernet port on the MDS 9000 switch; however, in this example, it is not needed because both the initiator and the Gigabit Ethernet port are on the same subnet.

Step 4 Confirm connectivity by pinging the Gigabit Ethernet port on the MDS 9509 switch. Similarly, ping the host NIC from the MDS 9509 switch Gigabit Ethernet port.

Note It is critical to check the connectivity between the host NIC card and the Gigabit Ethernet port on the MDS IPS blade before proceeding further. A successful ping test should be sufficient.

Step 5 Enable the iSCSI interface on the MDS 9509 switch.

The iSCSI interface 8/1 is the same port as the Gigabit Ethernet interface. At the same time you enable it, you can perform additional iSCSI related, TCP tuning. The following example uses the default values.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface iscsi 8/1
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 6 Configure the initiator on the MDS 9509 switch.

There are multiple ways to configure the initiator: using the IP address of the initiator, using the initiator as a proxy initiator, or using an IQN (iSCSI Qualified Name). This example uses the IQN. Most iSCSI drivers of clients can automatically assign an IQN to the host. The IQN name must be at least16 characters long. You can also manually assign the IQN name, but you must make sure that the IQN name is unique. The Windows 2000 server used in the Microsoft driver preconfigures the IQN name at the time of install.

Figure 8-2 show iSCSI initiator properties on the Windows 2000 server. In the **Initiator Settings** tab, you can see the IQN name for the host. The driver assigned initiator node name is iqn.1991-05.com.microsoft:sjc7-pc-1. For Linux systems, you can find this information in the /etc/initiatorname.iscsi file.

Figure 8-2 iSCSI Initiator Properties

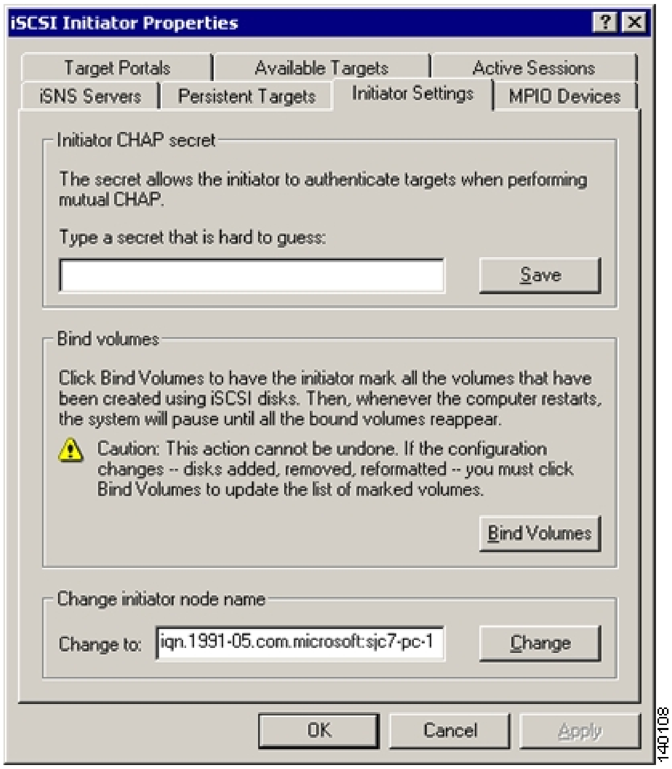

Initiator CHAP secret

The secret allows the initiator to authenticate targets when performing mutual CHAP.

Type a secret that is hard to guess:

Bind volumes

Click Bind Volumes to have the initiator mark all the volumes that have been created using iSCSI disks. Then, whenever the computer restarts, the system will pause until all the bound volumes reappear.

⚠ Caution: This action cannot be undone. If the configuration changes -- disks added, removed, reformatted -- you must click Bind Volumes to update the list of marked volumes.

Change initiator node name

Change to: iqn.1991-05.com.microsoft:sjc7-pc-1

To configure the iSCSI initiator on the MDS 9509 switch:

a. Manually specify the pWWN.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi initiator name iqn.1991-05.com.microsoft:sjc7-pc-1
sjc7-9509-6(config-(iscsi-init))# static pWWN 2105000dec022d82 <-- manually assigned
sjc7-9509-6(config-(iscsi-init))# vsan 3003 <-- Must be a member in the Targets VSAN
sjc7-9509-6(config-(iscsi-init))# ^Z
sjc7-9509-6#
```

b. Allow the MDS 9509 switch to determine the pWWN (preferred method).

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi initiator name iqn.1991-05.com.microsoft:sjc7-pc-1
sjc7-9509-6(config-(iscsi-init))# static pWWN system-assign 1 <-- system assigned
sjc7-9509-6(config-(iscsi-init))# vsan 3003 <-- Must be a member in the Targets VSAN
sjc7-9509-6(config-(iscsi-init))# ^Z
sjc7-9509-6#
```

The example uses the IQN assigned by the driver. If you need to change the name, make sure it is unique and at least 16 characters long. You can also optionally assign a pWWN to the initiator. The pWWN can be statically assigned by the administrator (as shown) or the system can automatically assign one. The initiator can be part of multiple VSANs. To access the target, the initiator has to be a member of the target's VSAN. In the previous example, the target belongs to VSAN 3003.

Note Alternatively, you can use the IP address of the iSCSI initiator for the configuration. Assigning a static pWWN is also optional. While configuring zoning for an iSCSI interface, you can use its IP address in place of its pWWN or the IQN name.

This is an example of the using the IP address for the iSCSI configuration.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi initiator ipaddress 10.1.45.2
sjc7-9509-6(config-(iscsi-init))# static pWWN system-assign 1 <-- system assigned
sjc7-9509-6(config-(iscsi-init))# vsan 3003 <-- Must be a member in the Targets VSAN
sjc7-9509-6(config-(iscsi-init))# ^Z
sjc7-9509-6#
```

Tip For the iSCSI initiator to communicate with a target port, the iSCSI initiator must be a member of the target port's VSAN. iSCSI initiators can be members of multiple VSANs.

Step 7 Configure a virtual target on the MDS 9509 switch.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi virtual-target name sym2127-fa4ba-sjc7
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-(iscsi-tgt))# ^Z
sjc7-9509-6#
```

Step 8 Assign a name to the virtual target. The name must be at least 16 characters long.

Step 9 To configure a virtual target, assign the pWWN of the storage port to the virtual target.

Step 10 Allow the initiator to communicate with the virtual target. Then the virtual target can permit all initiators or selected initiators to communicate with it, as the following example shows.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi virtual-target name sym2127-fa4ba-sjc7
sjc7-9509-6(config-(iscsi-tgt))#initiator iqn.1991-05.com.microsoft:sjc7-pc-1 permit
sjc7-9509-6(config-(iscsi-tgt))# ^Z
sjc7-9509-6#
```

In this example, the initiator iqn.1991-05.com.microsoft:sjc7-pc-1 can communicate with the virtual target sym2127-fa4ba-sjc7.

Step 11 Create a zone with the iSCSI initiator and the virtual target as members, so that the initiator and the target can communicate. You can create the zone with either the IQN name, the IP address, or the pWWN that was assigned to the initiator.

a. Zoning with the pWWN of the target and initiator:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# zone name Z_iscsi_sjc7-pc-1 vsan 3003
sjc7-9509-6(config-zone)# member pwwn 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-zone)# member pwwn 21:05:00:0d:ec:02:2d:82
sjc7-9509-6(config-zone)# ^Z
sjc7-9509-6#
```

b. Zoning with the pWWN of the virtual target and the IQN of the iSCSI initiator:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config-zone)# member pwwn 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-zone)# member symbolic-nodename iqn.1991-05.com.microsoft:sjc7-pc-1
sjc7-9509-6(config-zone)# ^Z
sjc7-9509-6#
```

Step 12 Activate the zone set and make the zone a member of the active zone set in VSAN 3003.

```
sjc7-9509-6(config)# zoneset name ZS_iSCSI vsan 3003
ssjc7-9509-6(config-zoneset)# member Z_iscsi_sjc7-pc-1
sjc7-9509-6(config-zoneset)# exit
sjc7-9509-6(config)# zoneset activate name ZS_iSCSI vsan 3003
Zoneset activation initiated. check zone status
sjc7-9509-6(config-zone)# ^Z
```

Step 13 Check the active zone set to see if they are both the target and initiator are active.

```
sjc7-9509-6# sh zoneset active vsan 3003
zoneset name ZS_iSCSI vsan 3003
zone name Z_iscsi_sjc7-pc-1 vsan 3003
  * fcid 0xd90002 [symbolic-nodename iqn.1991-05.com.microsoft:sjc7-pc-1]
  * fcid 0xd90000 [pwwn 50:06:04:82:bf:d1:db:d3]
sjc7-9509-6#
```

The iSCSI configuration on the MDS 9509 switch is complete.

Configuring the iSCSI Host

You can configure iSCSI clients on both Windows and Linux systems.

iSCSI Clients on Windows

This section describes a Microsoft Windows iSCSI driver 1.05a configuration.

To configure an iSCSI client on a Microsoft Windows 2000 system, follow these steps:

Step 1 In the iSCSI Initiator Properties window, click the **Target Portals** tab.

Step 2 Click **Add** to add a target portal on the Windows 2000 host, as shown in Figure 8-3. You see the **Add Target Portal** dialog box.

Figure 8-3 iSCSI Add Target Portals

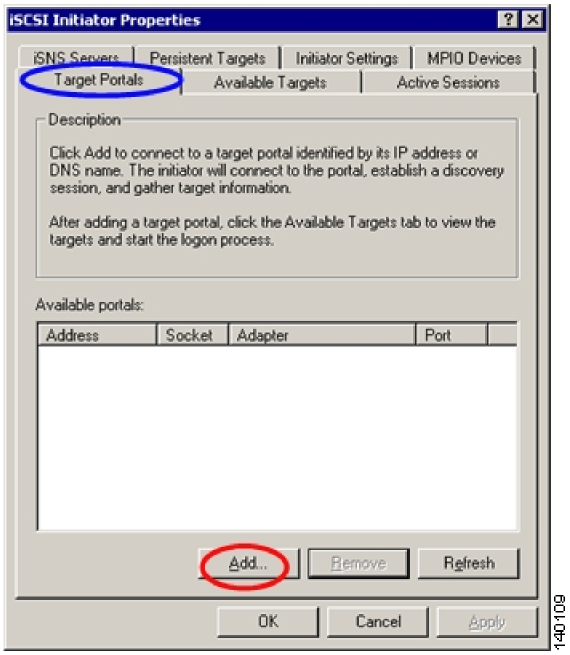

Click Add to connect to a target portal identified by its IP address or DNS name. The initiator will connect to the portal, establish a discovery session, and gather target information.

After adding a target portal, click the Available Targets tab to view the targets and start the logon process.

Step 3 In the **Add Target Portal** dialog box, enter the IP address assigned to the Gigabit Ethernet port on the MDS 9509 switch. In this case it is 10.1.45.1. See Figure 8-4.

Figure 8-4 Add Target Portal Window

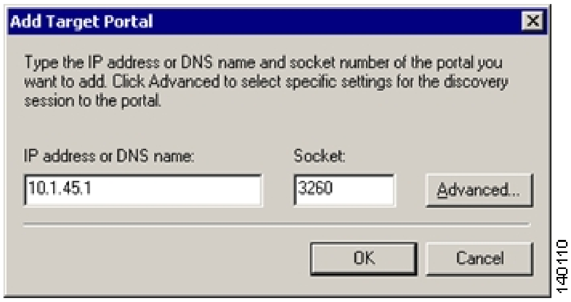

Type the IP address or DNS name and socket number of the portal you want to add. Click Advanced to select specific settings for the discovery session to the portal.

IP address or DNS name: 10.1.45.1

Socket: 3260

The target portal address appears in the list of target portals that the iSCSI client can log on to. (See Figure 8-5.)

Figure 8-5 iSCSI Interface Showing the Added Target

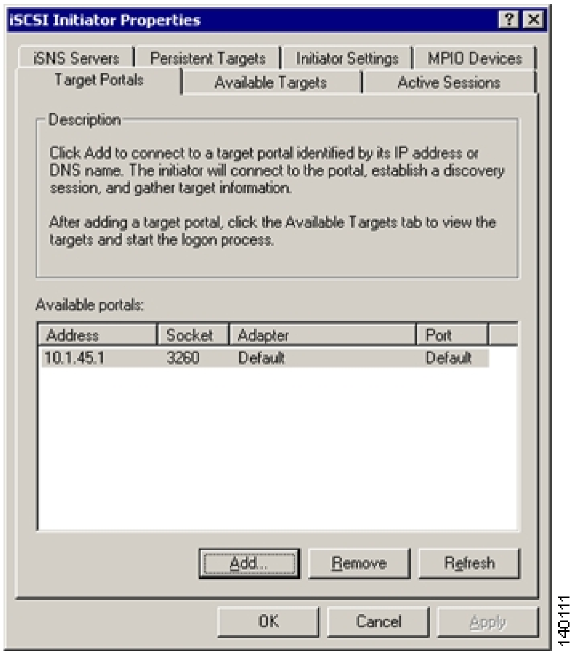

Description

Click Add to connect to a target portal identified by its IP address or DNS name. The initiator will connect to the portal, establish a discovery session, and gather target information.

After adding a target portal, click the Available Targets tab to view the targets and start the logon process.

Available portals:

Address	Socket	Adapter	Port
10.1.45.1	3260	Default	Default

Step 4 Log on to the target.

Step 5 Click the **Available Targets** tab in the iSCSI client window. If the client can see the target then the virtual configuration and zone of this initiator should be listed. (See Figure 8-6.) You may have to refresh the window to see it populated. If the target still does not appear, check the configuration on the switch and the iSCSI driver subsystem on the host.

Figure 8-6 iSCSI Available Targets View

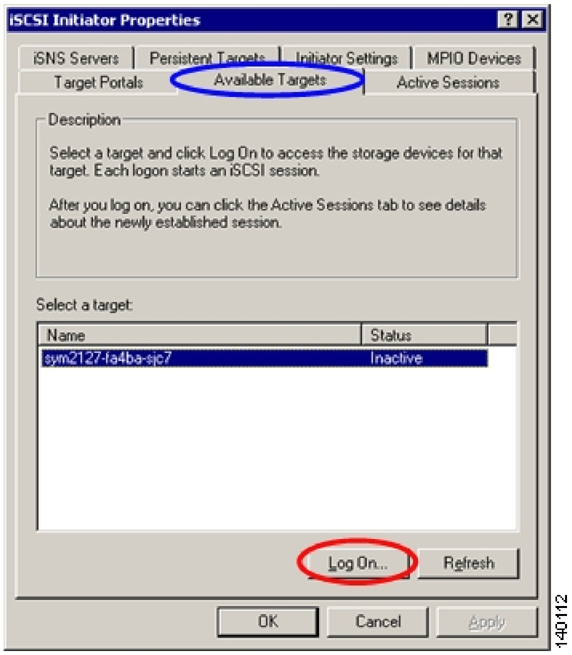

iSCSI Initiator Properties

Description

Select a target and click Log On to access the storage devices for that target. Each logon starts an iSCSI session.

After you log on, you can click the Active Sessions tab to see details about the newly established session.

Select a target:

Name	Status
sym2127-fa4ba-sjc7	Inactive

As Figure 8-6 shows, the status of the target is inactive.

Step 6 Click the target name and click **Log On**. You see the iSCSI **Log On to Target** dialog box shown in Figure 8-7. There are options to have the system automatically log on to the target and to enable multi-path if multi-pathing software is installed on the server.

Figure 8-7 Log On to Target Dialogue

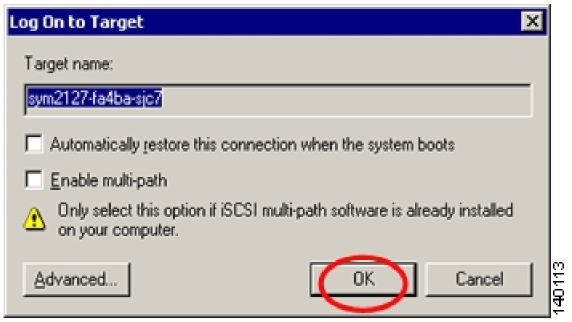

Log On to Target

Target name:

sym2127-fa4ba-sjc7

□ Automatically restore this connection when the system boots

□ Enable multi-path

⚠ Only select this option if iSCSI multi-path software is already installed on your computer.

Step 7 Click **OK** to start the iSCSI logon and storage discovery process for the target listed in the window. An iSCSI initiator can see multiple targets. The status of the target changes to Connected in the Available Targets tab once the initiator successfully logs in to the target. (See Figure 8-8.)

At this stage the host should be able to see the storage made available (directly assigned or made available through LUN masking or otherwise) to the iSCSI initiator.

Figure 8-8 Successful iSCSI Logon

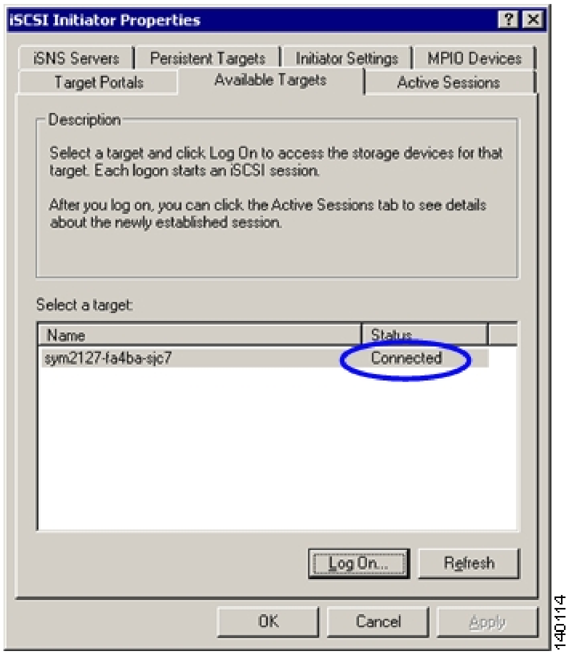

Description

Select a target and click Log On to access the storage devices for that target. Each logon starts an iSCSI session.

After you log on, you can click the Active Sessions tab to see details about the newly established session.

Select a target:

Name	Status
sym2127-fa4ba-sjc7	Connected

The **Active Sessions** tab and the Details button on that tab shows the LUNs visible to the initiator.

Step 8 From the Microsoft Windows **Start** button, choose **Settings > Control Panel > Administrative Tools > Computer Management > Disk Management** to scan for new disks and initialize them for use on the system.

iSCSI Clients on Linux

This section discusses the configuration of a Linux client to enable iSCSI on a Linux server. To enable iSCSI to work on a Linux host, you need to download the iSCSI driver and configure it. The iSCSI driver is available from this URL: http://sourceforge.net/project/showfiles.php?group_id=26396&release_id=177564. In some instances, you may need to compile the driver.

After you install the driver on the Linux host, you can configure it by editing the iscsi.conf file, which is typically in the /etc directory on a Linux server. As in the case of the Windows configuration, you need to add the target portal and configure the iSCSI initiator on an MDS 9000 switch. You also need to restart the iSCSI initiator on the Linux host and log on to the target to see the LUNs.

As in the case of the Windows iSCSI driver, the iSCSI initiator's name is auto-generated by the driver subsystem. The auto generated name is found in the /etc/initiatorname.iscsi file.

To configure the iSCSI initiator on a Linux host system, follow these steps:

Step 1 Edit the /etc/iscsi.conf file to enable iSCSI to work on the Linux server by adding the address of the target portal, which in this case is 10.1.45.1, in the DiscoverAddress Settings section.

```
# DiscoveryAddress Settings
# -------------------------
# Add "DiscoveryAddress=xxx" entries for each iSCSI router instance.
# The driver will attempt to discover iSCSI targets at that address
# and make as many targets as possible available for use.
# 'xxx' can be an IP address or a hostname.  A TCP port number can be
# specified by appending a colon and the port number to the address.
# All entries have to start in column one and must not contain any
# whitespace.
#
# Example:
#
DiscoveryAddress=10.1.45.1
```

These changes configure basic iSCSI on the Linux server.

Step 2 Restart the iSCSI process to force the iSCSI client to log on to the target and discover the LUNs. You do this using the rc script, S25iscsi, which is located in /etc/rc3.d directory. The script also has a status option to check the status of the iscsi processes on the system.

```
[root@sjc7-pc-6 rc3.d]# /etc/rc3.d/S25iscsi restart
Stopping iSCSI: sync umount sync iscsid iscsi
Starting iSCSI: iscsi iscsid

[root@sjc7-pc-6 rc3.d]# /etc/rc3.d/S25iscsi status
iSCSI driver is loaded
[root@sjc7-pc-6 rc3.d]#
```

The iSCSI initiator should log on to the target and discover the LUNs that are assigned to it.

Configuring iSCSI in Proxy Initiator Mode

The following procedure presents the Proxy Initiator Mode configuration for iSCSI on an Cisco MDS 9000 switch. This method is the preferred mode for configuring a large number of iSCSI clients to work with the switch.

In Proxy Initiator Mode, the one Fibre Channel initiator is used for all iSCSI clients that access the MDS 9000 switch via the same iSCSI interface (iscsi3/3, for example). The initiators use the pWWN assigned to the iSCSI interface. The iSCSI interface that an iSCSI client logs in to is configured in the client and must be permitted by the virtual target configured for that initiator.

Proxy Initiator Mode is advantageous over Transparent Mode when the configuration requires multiple iSCSI initiators to access the same Fibre Channel target. For example, if 20 iSCSI initiators need to communicate with a Fibre Channel disk, Transparent Mode requires that 20 iSCSI initiators and 20 zones need to be created. In addition, array based LUN masking has to be updated for all 20 initiator instances.

On the other hand, Proxy Initiator Mode is far easier to manage as it allows for centralized management of the iSCSI configuration, as all iSCSI clients accessing the same switch interface use a single Fibre Channel initiator.

First, you assign a pWWN to the iSCSI interface. Then you zone this one pWWN with the Fibre Channel target so that the proxy initiator can see the LUNs presented by the virtual target. All the LUN masking and zoning need to be performed only with the proxy initiator. As you add new hosts (iSCSI clients), you only need to expose them to the LUNs that they need to see. No new zones or modifications to the array's LUN masking are needed.

A typical practice is to create a virtual target for each host and configure the virtual target in such a way that it only exposes the required LUNs to the iSCSI initiator.

The proxy initiator is not restricted to a single VSAN. As iSCSI clients are configured and given access to different VSANs, the proxy initiator is created in the new VSAN. Therefore, the maximum number of initiators that need to be zoned is the number of proxy initiators that have iSCSI clients in a VSAN. This is far less than under Transparent Mode, whereby a Fibre Channel initiator is created for *every* iSCSI client.

The topology used for the iSCSI proxy initiator recipe appears in Figure 8-9. It has two Windows 2000 hosts on two different subnets. The first host's iSCSI interface is on 10.1.45.0 network and the second interface is on 10.1.46.0 network.

Figure 8-9 iSCSI Proxy Initiator Topology

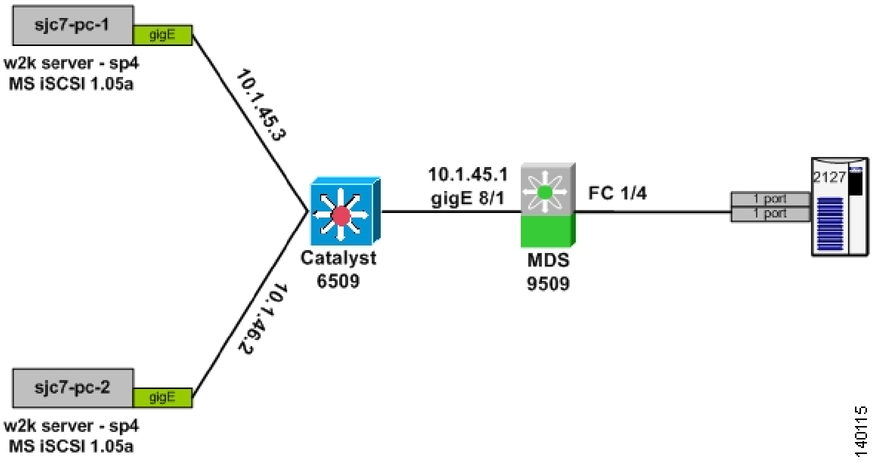

To create a proxy initiator configuration, follow these steps:

Step 1 Configure the Gigabit Ethernet interface on the MDS 9509 switch by assigning it an IP address and subnet mask. The address allows the Gigabit Ethernet interface to communicate with the network.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface gigabitethernet 8/1
sjc7-9509-6(config-if)# ip address 10.1.45.1 255.255.255.0
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 2 Configure the IP route if required

The IP route should be configured to communicate with the initiator if the initiator is in another subnet. In the preceding step, both the initiator and the iSCSI interface on the MDS 9509 switch are in the same subnet. Therefore, the configuration of an explicit route on the MDS 9509 switch is not required.

The following example shows the syntax for configuring a route. The initiator is in the 10.1.46.0 subnet, so you need to add a route to reach that subnet from the MDS 9509 switch.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# ip route 10.1.46.0 255.255.255.0 10.1.45.2
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 3 Enable and configure the iSCSI interface on the MDS 9509 switch.

The iSCSI interface 8/1 is the same port as the Gigabit Ethernet interface. At the same time as you enable the iSCSI interface, you can perform additional iSCSI related TCP tuning. This example uses the default values. Use the switch port command to enable Proxy Initiator Mode for the iSCSI interface 8/1. It assigns pWWN to the interface as shown.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface iscsi 8/1
sjc7-9509-6(config-if)# switchport proxy-initiator
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 4 Configure the iSCSI interface to access multiple VSANs. This allows the iSCSI interface to communicate with the virtual target to see the LUNs. The commands below add the iSCSI interface 8/1 into VSAN 1. All clients not explicitly configured for a VSAN belong to VSAN 1.

Tip Using VSAN 1 as a default, non-production VSAN, prevents unconfigured devices from being able to access devices in a production VSAN. In this case, the production VSAN is 3003.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi interface vsan-membership
sjc7-9509-6(config)# vsan database
sjc7-9509-6(config-vsan-db)# vsan 1 interface iscsi 8/1
sjc7-9509-6(config-vsan-db)#^Z
sjc7-9509-6#
```

Step 5 Configure the iSCSI initiator.

With Proxy Initiator Mode, the proxy initiator has a default VSAN for all iSCSI clients that are not explicitly configured for a particular VSAN. The default value is VSAN 1. Having the default VSAN allows a single iSCSI client to access Fibre Channel targets in multiple VSANs. To add multiple VSANs, specify more VSANs in the command.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi initiator ipaddress 10.1.45.2
sjc7-9509-6(config-(iscsi-init))# static pWWN system-assign 1 <-- system assigned
sjc7-9509-6(config-(iscsi-init))# vsan 3003 <-- Must be a member in the Targets VSAN
sjc7-9509-6(config-(iscsi-init))# ^Z
sjc7-9509-6#

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi initiator ipaddress 10.1.45.3
sjc7-9509-6(config-(iscsi-init))# static pWWN system-assign 1 <-- system assigned
sjc7-9509-6(config-(iscsi-init))# vsan 3003 <-- Must be a member in the Targets VSAN
sjc7-9509-6(config-(iscsi-init))# ^Z
sjc7-9509-6#
```

Note You must use the **iscsi interface vsan-membership** command to make the iSCSI interface part of a VSAN.

Step 6 Configure a virtual target on the MDS 9509 switch by assigning it a name that is at least 16 characters long and assigning the pWWN of the storage port to the virtual port.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi virtual-target name sym2127-fa4ba-sjc7
ssjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-(iscsi-tgt))# ^Z
sjc7-9509-6#
```

Step 7 Create a zone with the initiator and the virtual target by creating a zone with the iSCSI proxy interface and the virtual targets as members. This zone enables the initiator and the target to communicate. You can create the zone with the interface iSCSI 8/1, with the IP address, or with the pWWN that is assigned to the initiator.

You can obtain the pWWN of the proxy initiator by viewing the iSCSI interface.

```
sjc7-9509-6# show int iscsi 3/3
iscsi8/1 is up
    Hardware is GigabitEthernet
    Port WWN is 20:89:00:0c:85:e9:d2:c0
    Admin port mode is ISCSI
    Port mode is ISCSI
    Port vsan is 1
    Speed is 1 Gbps
    iSCSI initiator is identified by name
    Number of iSCSI session: 0, Number of TCP connection: 0
    Configured TCP parameters
        Local Port is 3260
        PMTU discover is enabled, reset timeout is 3600 sec
        Keepalive-timeout is 60 sec
        Minimum-retransmit-time is 300 ms
        Max-retransmissions 4
        Sack is enabled
        QOS code point is 0
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 1000000 kbps
        Estimated round trip time is 1000 usec
        Send buffer size is 4096 KB
        Congestion window monitoring is enabled, burst size is 50 KB
        Configured maximum jitter is 500 us
    Forwarding mode: pass-thru
    TMF Queueing Mode : disabled
    Proxy Initiator Mode : enabled
        nWWN is 21:04:00:0d:ec:02:2d:82 (system-assigned)
        pWWN is 21:05:00:0d:ec:02:2d:82 (system-assigned)
```

a. Zoning with the pWWN of the target and proxy initiator:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# zone name Z_iscsi_sjc7-pc-1 vsan 3003
sjc7-9509-6(config-zone)# member pwwn 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-zone)# member pwwn 21:05:00:0d:ec:02:2d:82
sjc7-9509-6(config-zone)# ^Z
sjc7-9509-6#
```

b. Zoning with the pWWN of the virtual target and the IP address of the iSCSI interface:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# zone name Z_iscsi_sjc7-pc-1 vsan 3003
sjc7-9509-6(config-zone)# member pwwn 50:06:04:82:bf:d1:db:d3
sjc7-9509-6(config-zone)# member member ip-address 10.1.45.1 255.255.255.0
sjc7-9509-6(config-zone)# ^Z
sjc7-9509-6#
```

Step 8 Activate the zone set by making it a member of the active zone set in VSAN 3003.

```
sjc7-9509-6(config)# zoneset name ZS_iSCSI vsan 3003
sjc7-9509-6(config-zoneset)# member Z_iscsi_sjc7-pc-1
sjc7-9509-6(config-zoneset)# exit

sjc7-9509-6(config)# zoneset activate name ZS_iSCSI vsan 3003
Zoneset activation initiated. check zone status
sjc7-9509-6(config-zone)# ^Z

sjc7-9509-6# show zoneset active vsan 3003
zoneset name ZS_iscsi vsan 3003
  zone name Z_iscsi_proxy_8-1 vsan 3003
  * fcid 0xd90002 [pwwn 21:05:00:0d:ec:02:2d:82]
  * fcid 0xd90000 [pwwn 50:06:04:82:bf:d1:db:d3]
sjc7-9509-6#
```

Step 9 Configure a virtual target for each initiator and configure LUN masking for the initiator.

Once the zone is successfully activated, the LUNs that are made available on the storage port should be visible to the iSCSI interface. As this interface could be a proxy iSCSI interface for many iSCSI initiators, some form of LUN security must be enabled. To achieve LUN security, create a virtual target with access to specific LUNs for each initiator. In this example, the iSCSI interface can see 10 LUNs (LUN 11 to LUN 20 in decimal).

a. This configuration allows the host sjc7-pc-1 to see LUNs 11 - 14 (decimal) on the array.

```
sjc7-9509-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi virtual-target name sym2127-fa4ba-sjc7-1
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun b iscsi-lun 1
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun b iscsi-lun 2
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun d iscsi-lun 3
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun e iscsi-lun 4
sjc7-9509-6(config-(iscsi-tgt))# initiator ip address 10.1.45.3 permit
sjc7-9509-6(config-(iscsi-tgt))# end
sjc7-9509-6#
```

b. Similarly for the host jc7-pc-2, the following configuration allows it to see LUNs 16 - 20 (decimal).

```
sjc7-9509-6# config t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# iscsi virtual-target name sym2127-fa4ba-sjc7-2
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun 10 iscsi-lun 1
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun 11 iscsi-lun 2
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun 12 iscsi-lun 3
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun 13 iscsi-lun 4
sjc7-9509-6(config-(iscsi-tgt))# pwwN 50:06:04:82:bf:d1:db:d3 fc-lun 14 iscsi-lun 5
sjc7-9509-6(config-(iscsi-tgt))# initiator ip address 10.1.46.2 permit
sjc7-9509-6(config-(iscsi-tgt))# end
sjc7-9509-6
```

After the above configuration changes, both the hosts should be able to see the LUNs allocated to them through the virtual target created for each. There is no need to create additional zones when new iSCSI clients are added and they only need to see the target already in the zone. If the iSCSI clients need access to additional targets, the iSCSI interface needs to be zoned to the other targets as shown in Step 6 and Step 7.

The client side iSCSI configuration is the same as that described in Configuring the iSCSI Host, page 8-6.

CHAPTER 9

FCIP

FCIP (Fibre Channel over IP) is an IETF standards based protocol for connecting Fibre Channel SANs over IP based networks. FCIP encapsulates the FCP frames in a TCP/IP packet which is then sent across an IP network. FCIP can interconnect geographically dispersed SANs using the IP backbone network. In short, FCIP creates an FC tunnel over an existing IP network. The MDS 9200 and MDS 9500 series switches support FCIP, using the IPS-8, IPS-4 and the 14+2 blades.

This chapter explains how to enable and configure FCIP and includes the following sections:

- Enabling FCIP, page 9-1
- Configuring FCIP, page 9-1
- Tuning FCIP, page 9-5

Enabling FCIP

You must enable FCIP before attempting to configure it on the switch.

Warning **If you do not run the FCIP enable command, further FCIP configuration is not possible. This command enables additional FCIP configuration options in the CLI.**

```
MDS1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
MDS1(config)# FCIP enable
MDS1(config)#^Z
MDS1#
```

Configuring FCIP

The recipe below shows the configuration of FCIP on the Cisco MDS 9000 switch. Figure 9-1 shows the topology.

Figure 9-1 FCIP Topology

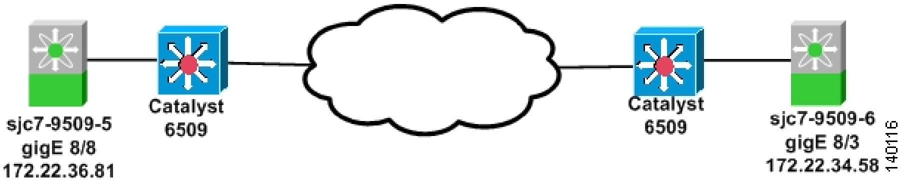

The topology consists of two MDS 9509 switches, each with one IPS-8 blade. The Gigabit Ethernet port 8/8 on the switch sjc7-9509-5 is connected to a Catalyst 6509. The interface has an IP address of 172.22.36.81, with a subnet mask of 255.255.254.0, and the gateway address of 172.22.36.1.Another switch, sjc7-9509-6, has the Gigabit Ethernet port 8/3 also connected to a Catalyst 6509. This interface has an IP address of172.22.34.58, with a subnet of 255.255.254.0, and the gateway address of 172.22.34.1. In the recipe, the FCIP tunnel is established between the switch sjc7-9509-5 (GE 8/8) and sjc7-9509-6 (GE 8/3).

Warning **The IP address for the ports on the IPS blade must be on a different subnet and Ethernet segment than the management interface. The different segment can be achieved either though physical separation or using VLANs. Having different segments is critical to get FCIP working on the MDS 9000 Family switch.**

To make FCIP operational between the switches, follow these steps:

Step 1 Configure the Gigabit Ethernet port on the MDS 9000 switch sjc7-9509-5 by assigning an IP address and a subnet mask. The address and subnet mask allow the Gigabit Ethernet interface to communicate with the network.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface gigabitethernet 8/3
sjc7-9509-6(config-if)# ip address 172.22.34.58 255.255.254.0
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-5#
```

Step 2 Configure IP routes so that the two Gigabit Ethernet interfaces on switches sjc7-9509-5 and sjc7-9509-6 can communicate with each other.

In the preceding step, the Gigabit Ethernet ports are in two different subnets. As a result, explicit routes are required to allow the two switches to communicate with each other.

Note **The recommended best practice for configuring an IP route is to have a host route to each of the two Gigabit Ethernet interfaces. In other words, you should allow the two Gigabit Ethernet interfaces to communicate only with each other. You can accomplish this by designating a subnet mask of 255.255.255.255 when you add the route.**

Step 3 Use the following syntax to configure a route to allow the Gigabit Ethernet port 8/8 on the switch sjc7-9509-5 to communicate with the Gigabit Ethernet port 8/3 on switch sjc7-9509-6:

```
sjc7-9509-5# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# ip route 172.22.34.58 255.255.255.255 172.22.36.1 interface
gigabitethernet 8/8
sjc7-9509-5(config)# ^Z
```

```
sjc7-9509-5#
```

In this configuration, you use the gateway 172.22.36.1 and interface GE 8/8 on switch sjc7-9509-5 to reach 172.22.34.58.

Step 4 Similarly, use the following syntax to configure a route for GE 8/3 on switch sjc7-9509-6 to communicate with the port GE 8/8 on switch sjc7-9509-5:

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# ip route 172.22.36.81 255.255.255.255 172.22.34.1 interface
gigabitethernet 8/3
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

In this configuration, you use the gateway 172.22.34.1 and interface GE 8/3 on switch sjc7-9509-6 to reach 172.22.36.81.

Step 5 Ping the Gigabit Ethernet interfaces to ensure that the Gigabit Ethernet ports can communicate with each other.

a. From the switch prompt on sjc7-9509-5, ping the IP address of the Gigabit Ethernet interface 8/3 on switch sjc7-9509-6.

b. From the switch prompt on sjc7-9509-6, ping the IP address of the Gigabit Ethernet interface 8/8 on switch sjc7-9509-5.

```
sjc7-9509-5# ping 172.22.34.58
PING 172.22.34.58 (172.22.34.58) 56(84) bytes of data.
64 bytes from 172.22.34.58: icmp_seq=1 ttl=254 time=1.04 ms
64 bytes from 172.22.34.58: icmp_seq=2 ttl=254 time=0.624 ms
64 bytes from 172.22.34.58: icmp_seq=3 ttl=254 time=0.678 ms
64 bytes from 172.22.34.58: icmp_seq=4 ttl=254 time=0.580 ms

--- 172.22.34.58 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3008ms
rtt min/avg/max/mdev = 0.580/0.732/1.046/0.184 ms
sjc7-9509-5#

sjc7-9509-6# ping 172.22.36.81
PING 172.22.36.81 (172.22.36.81): 56 data bytes
64 bytes from 172.22.36.81: icmp_seq=0 ttl=254 time=0.7 ms
64 bytes from 172.22.36.81: icmp_seq=1 ttl=254 time=0.6 ms
64 bytes from 172.22.36.81: icmp_seq=2 ttl=254 time=0.5 ms
64 bytes from 172.22.36.81: icmp_seq=3 ttl=254 time=0.6 ms

--- 172.22.36.81 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 0.5/0.6/0.7 ms
sjc7-9509-6#
```

Note **It is critical to check the connectivity between the two Gigabit Ethernet ports on each MDS 9000 Family switch IPS blade before proceeding further. A successful ping test should be sufficient to verify connectivity.**

Step 6 Measure the round-trip time between the two Gigabit Ethernet interfaces because you will need the value later in the configuration.

```
sjc7-9509-5# ips measure-rtt 172.22.34.58 interface gigabitethernet 8/8
Round trip time is 691 micro seconds (0.69 milli seconds)
sjc7-9509-5#

sjc7-9509-6# ips measure-rtt 172.22.36.81 interface gigabitethernet 8/3
Round trip time is 743 micro seconds (0.74 milli seconds)
```

```
sjc7-9509-6#
```

Step 7 Configure the FCIP profile on both the switches by creating the FCIP profile.

The profile defines the characteristics for the FCIP tunnel that you set up. The measured round-trip time is needed during the profile configuration. Round the value to the nearest whole number. In this case, the 690 microseconds is rounded to 1 millisecond. The IP address in the configuration that follows is the IP address assigned to the Gigabit Ethernet interface on the switch.

Tip We recommend that the FCIP profile and FCIP interface numbers at both ends of a FCIP tunnel be the same for easy management.

```
sjc7-9509-5# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# fcip profile 10
sjc7-9509-5(config-profile)# ip address 172.22.36.81
sjc7-9509-5(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth-mbps 200
round-trip-time-ms 1
sjc7-9509-5(config-if)# ^Z
sjc7-9509-5#

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# fcip profile 10
sjc7-9509-6(config-profile)# ip address 172.22.34.58
sjc7-9509-6(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth-mbps 200
round-trip-time-ms 1
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

Step 8 Configure the FCIP interface on both switches.

In the FCIP interface configuration that follows, the profile used and the peer information (remote Gigabit Ethernet IP address) are specified. Compression and write acceleration can also be configured.

```
sjc7-9509-5# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# interface fcip 1
sjc7-9509-5(config-if)# use-profile 10
sjc7-9509-5(config-if)# peer-info ipaddr 172.22.34.58
sjc7-9509-5(config-if)# no shut
sjc7-9509-5(config-if)# ^Z
sjc7-9509-5#

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface fcip 1
sjc7-9509-6(config-if)# use-profile 10
sjc7-9509-6(config-if)# peer-info ipaddr 172.22.36.81
sjc7-9509-6(config-if)# no shut
sjc7-9509-6(config-if)# ^Z
sjc7-9509-6#
```

The FCIP tunnel should be up and running now.

Step 9 Use the **show interface fcip 1 brief** command to show the status of the FCIP link between the two switches.

```
sjc7-9509-5#  show interface fcip 1 brief

-------------------------------------------------------------------------------
Interface Vsan Admin Admin  Status       Oper Profile    Eth Int    Port-channel
               Mode  Trunk               Mode
                     Mode
-------------------------------------------------------------------------------
fcip1     1    auto  on    trunking      TE    10 GigabitEthernet8/8  --
sjc7-9509-5#

sjc7-9509-6#  show interface fcip 1 brief

-------------------------------------------------------------------------------
Interface Vsan Admin Admin  Status       Oper Profile    Eth Int    Port-channel
               Mode  Trunk               Mode
                     Mode
-------------------------------------------------------------------------------
fcip1     1    auto  on    trunking      TE    10 GigabitEthernet8/3 --
sjc7-9509-6#
```

Tuning FCIP

Implementing an FCIP tunnel is more than just creating the FCIP link and modifying the VSAN allowed list. To achieve the greatest efficiency from the link, some parameters that are specific to the underlying connection may need to be tuned. Therefore, if the FCIP link is running over a slow, 1 MB connection, the FCIP link should be tuned differently than one running over a low-latency, 1 GB connection. This section provides insight into what parameters can be used and how to use them for achieving increased efficiency and utilization from the FCIP connection.

Note Your results may vary due to network conditions (existing utilization) and storage array and host type.

Figure 9-2 shows the topology that is used throughout this chapter:

Figure 9-2 FCIP Topology

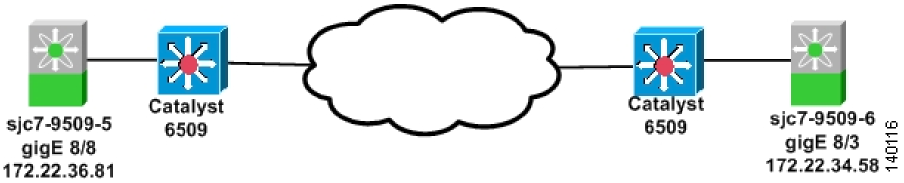

TCP Tuning: Latency and Available Bandwidth

The latency of the link is the amount of time that it takes a packet to go from one end of the FCIP link to the other. Latency can be due to many factors, including distance and the number of devices that the packet must traverse. Even the fastest routers and switches incur some amount of latency.

Latency cannot be eliminated; however protocols can be tuned and MDS features such as FCIP Write Acceleration (see FCIP Write Acceleration, page 9-7) can be enabled to eliminate the effects of it. These features can be modified in the FCIP profile.

Available bandwidth is the amount of bandwidth that the FCIP link can use on the network. You need to define a maximum and a minimum value for the FCIP link to use in the FCIP profile.

- The maximum available bandwidth value is the maximum amount of bandwidth that the FCIP link should use on the network.
- The minimum available bandwidth value is used as a guideline for the minimum value. Of course, if there are serious problems on the network (such as dropped packets or congestion), the link will go slower than the minimum value. We recommend that the minimum value be set to what the minimum acceptable bandwidth that the application (EMC SRDF, IBM PPRC, and so on) requires.

Tip For dedicated links, set the minimum and maximum available bandwidth values to be the same.

Table 9-1 contains some common WAN links and their speeds. These circuits are most often used as the underlying network for an FCIP link. For example, the underlying network may be a OC3, but you may only be able to use 100 MB of that link.

Tip When deploying FCIP, you should always involve the LAN and WAN teams to find out about the connectivity that they are providing you. If there are performance issues, they can often help you troubleshoot them from the network's standpoint. Involve them earlier rather than later.

Table 9-1 Common WAN Circuit Speeds

Circuit Name	Speed
T1	1.544 Mbps
T3	43.232 Mbps
OC-3	155 Mbps
OC-12	622 Mbps
OC-48	2.5 Gbps
OC-192	9.6 Gbps

To specify the minimum and maximum available bandwidth and the network latency, perform the following procedure. For this example, we assume that the LAN/WAN links are a dedicated Gigabit Ethernet link.

Step 1 Measure the round-trip time between the two FCIP end points.

```
sjc7-9509-5# ips measure-rtt 172.22.34.58 interface gigabitethernet 8/8
Round trip time is 5691 micro seconds (5.69 milli seconds)

sjc7-9509-6# ips measure-rtt 172.22.36.81 interface gigabitethernet 8/3
Round trip time is 5743 micro seconds (5.74 milli seconds)
```

Step 2 Modify the profile by specifying the minimum and maximum bandwidth.

```
sjc7-9509-5# conf t
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# fcip profile 2
sjc7-9509-5(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth
-mbps 1000 round-trip-time-us 5691

sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# fcip profile 2
sjc7-9509-6(config-profile)# tcp max-bandwidth-mbps 1000 min-available-bandwidth
-mbps 1000 round-trip-time-us 5743
```

FCIP Write Acceleration

To alleviate the effects of latency, the IPS-8 and IPS-4 blades support software compression. The configuration is shown below.

Note Only in SAN-OS Release 2.0(1b) and higher should write acceleration be used with PortChannels.

Step 1 Enable write acceleration on switch sjc7-9509-5.

```
sjc7-9509-5# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-5(config)# interface fcip 1
sjc7-9509-5(config)# write-accelerator
sjc7-9509-5(config)# ip-compression mode high-throughput
sjc7-9509-5(config)# ^Z
sjc7-9509-5#
```

Step 2 Enable write acceleration on switch sjc7-9509-6.

```
sjc7-9509-6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
sjc7-9509-6(config)# interface fcip 1
sjc7-9509-6(config)# write-accelerator
sjc7-9509-6(config)# ip-compression mode high-throughput
sjc7-9509-6(config)# ^Z
sjc7-9509-6#
```

FCIP Compression

FCIP based compression should be enabled when the bandwidth requirements of the application (such as SRDF, TrueCopy, or PPRC) exceed the bandwidth that the LAN or WAN link can provide.

For SAN-OS versions prior to 2.0, software compression was available on the IP Services Module using two compression modes, a high compression ratio, and high throughput. For environments where the bandwidth speed of the WAN is relatively low (such as T3/DS 3, 45 Mbps and lower speeds), the high compression ratio option should be enabled. For WAN links up to OC-3, 155 Mbps, the high throughput option should be enabled.

In SAN-OS version 2.0, three new modes of compression were introduced, along with hardware based compression using the 14+2 module. These modules are summarized in Figure 9-3.

Note The numbers displayed in Figure 9-3 are derived using the Canterbury Corpus test suite. Your actual performance numbers may vary depending on device type and data pattern. However, the numbers shown here should be used as a guide to determine which compression mode is most appropriate.

Figure 9-3 Approximate Application Throughput with MDS 2.0 Compression

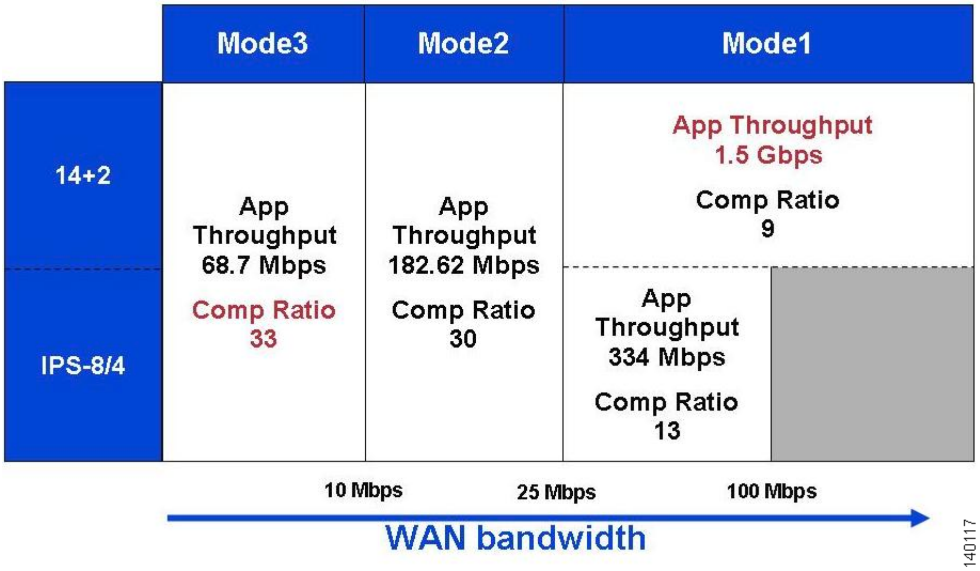

INDEX